THE COUNTERVAILING FORCES OF SELECTION AND BINDING IN VISION

by

ADAM CHRISTOPHER SNYDER

A dissertation submitted to the Graduate Faculty in Psychology in partial fulfillment of the requirements for the degree of Doctor of Philosophy, The City University of New York

2011

This manuscript has been read and accepted for the
Graduate Faculty in Psychology in satisfaction of the
dissertation requirement for the degree of Doctor of Philosophy.

Date	John J. Foxe, Ph.D. Chair of Examining Committee
Date	Maureen O’Connor, Ph.D. Executive Officer

Supervisory Committee:

Simon P. Kelly, Ph.D.

Jonathan B. Levitt, Ph.D.

Sophie Molholm, Ph.D.

Steven A. Hillyard, Ph.D.

Steven Yantis, Ph.D.

THE CITY UNIVERSITY OF NEW YORK

Abstract

THE COUNTERVAILING FORCES OF SELECTION AND BINDING IN VISION

by

Adam Christopher Snyder

Adviser: John J. Foxe, Ph.D.

Humans have limited cognitive resources to process the nearly limitless information available in the environment. Endogenous, or “top-down”, selective attention to basic visual features such as color or motion is a common strategy for biasing resources in favor of the most relevant information sources in a given context. Opposing this top-down separation of features is a “bottom-up” tendency to integrate, or bind, the various features that constitute objects. First we identify an electrophysiological signature of endogenous selective attention to basic visual features: alpha-band power increases, which have been shown to reflect suppression of potentially distracting information in several contexts. Next, we pitted the two processes of selective attention and binding against each other in a series of behavioral and electrophysiological experiments to test if top-down selective attention can overcome constitutive binding processes. Our results demonstrate that bottom-up binding processes can dominate top-down feature-based attention even when explicitly detrimental to task performance. A model for anatomical bases and temporal dynamics of the deployment of feature-based selective attention is proposed.

Dedication

To my wife, Bonnie.

Acknowledgments

Much gratitude is owed to my mentor, John Foxe, not only for the education, support and opportunities he provided me during my graduate studies, but also for bringing together a diverse group of students from wide-ranging academic backgrounds (yet united by a common curiosity and skepticism), which I feel is why our lab is such an enriching and successful scientific environment. Many thanks also to the senior researchers in the lab, Sophie Molholm and Pierfilipo De Sanctis, for their plentiful guidance. Thanks to the many post-doctoral fellows that educated me in the technical aspects of electrophysiology: John Butler, Ian Fiebelkorn, Hans-Peter Frey, Manuel Gomez-Ramirez, Simon Kelly, Ed Lalor, Manuel Mercier and Lars Ross. I would also like to thank my colleagues in the program, too numerous to list, for their help and friendship over the last four years. Very special thanks are due to Edel Flynn, who unswervingly solved every administrative challenge I faced in graduate school. I would also like to thank Sara Schildkraut, for help with data collection. Four of my pre-graduate instructors can be credited with cultivating my interest in science into a passion: two inspirational high school science teachers, Diana Prout and Cindy Taylor, and two college professors, Dan Sanes and Clay Curtis. Finally, I would like to thank my family for their eternal support, and for instilling in me the value of critical thinking.

Table of Contents

List of Tables

List of Figures

GENERAL INTRODUCTION

The ability to selectively attend to a specific subset of one's sensory realm is at the root of human cognition and negotiation of one’s environment. The alternative, whereby attention is not constrained to a manageable subset of the available information, does not allow for normal function, and may be partially involved in certain neural disorders such as the Autism spectrum of disorders (cf. Markram and Markram, 2010) and attention deficit disorder (Armstrong, Hayes and Martin, 2001). By necessity, successful task completion in daily life employs selective attention to one extent or another. By inference, the ability to attend is assumed to be a “capacity-limited” system and to some degree, humans have endogenous or executive control over what aspects of their environment are processed by this limited system from moment to moment (e.g. Broadbent, 1958, Kahneman, 1973; Schneider and Shiffrin, 1977, Shiffrin and Schneider, 1977; Neely, 1977; Schneider and Fisk, 1982). This endogenous form of attention can be distinguished from exogenous attention, which is involuntarily captured by properties of the environment. Exogenous attention can be captured by saliency cues such as an unexpected sound (Dalton and Lavie, 2004) or a sudden flash (Yantis and Jonides, 1984), or by cues such as biological relevance of an object category (e.g., snakes to a primate [Lipp and Waters, 2007] or cigarettes to a smoker [Versace et al., 2010]). This volume will describe a subtler involuntary capture of attention caused by the inherent bias in humans to perceive the world as composed of objects. **Chapter 1** will discuss how endogenous attention is oriented to specific features of the visual world, and **Chapters 2** and **3** will describe how this endogenous attention interacts with the automatic capture of attention to integrate disparate features to form coherent objects.

The Cuing Paradigm

In order to study the mechanisms of attention, we have employed the “cuing paradigm” to invoke endogenous biasing of the attentional system. This is an effective experimental means of ensuring invocation of a high level of endogenous control by the use of symbolic cues that actively instruct the subject in an unpredictable manner to switch attention between tasks. That is, a centrally presented cue such as an arrow, word or sound is employed (stimulus 1 – S1), which informs of an upcoming task or indicates a specific location, sensory modality or feature for subsequent analysis (stimulus 2 – S2). This manipulation is a nice experimental version of what is a fairly common way of apportioning attention in our environment. That is, humans spend a large part of their existence assigning attentive resources to their environment based on an internal instructional set.

Such endogenous cuing paradigms can be classified into two basic categories: “probabilistic” and “instructional”. In a probabilistic cuing paradigm, participants are instructed to respond to all target S2’s, including those that may occur at an uncued location. Critically, the cue indicates the correct location of the target on a high proportion of trials, which encourages the participants to utilize the cue’s information. The uncued location in this case cannot be considered a source of distracters, however, since participants still need to respond to targets there. In contrast, an instructional cuing paradigm requires the participants to respond to targets occurring at the cued location only, and to ignore the uncued location completely. In this way, the uncued location becomes simply a source of potentially distracting information, since stimuli occurring there will never warrant a response. On the other hand, since no response is ever made to such distracters, a physiological measure –rather than a behavioral measure –is needed to evaluate the processing thereof.

SELECTIVE ATTENTION

Two basic varieties of attentionally mediated, selective modulation of the neural circuitry involved in processing a particular stimulus array can be envisioned. The first, and by far the most often reported, involves the selective enhancement of neural responsiveness or efficacy in those cells responsible for processing the particular stimulus of interest –the stimulus, whether already present or merely anticipated, upon which a subject's attention is focused (e.g. Corbetta, Miezin, Dobmeyer, Shulman and Petersen, 1993; Haenny and Schiller, 1988; Luck, Chelazzi, Hillyard and Desimone, 1997; Moran and Desimone, 1985). The second variety of attentional modulation involves selective suppression of those neurons responsible for processing stimuli outside the focus of attention, so that competing or distracting information is not or *will not be* processed fully (e.g. Foxe, Simpson and Ahlfors, 1998; Moran and Desimone, 1985; Motter, 1993; Reynolds, Chelazzi and Desimone, 1999; Vanduffel, Tootell and Orban, 2000; Worden, Foxe, Wang and Simpson, 2000). The experiments described in this report emphasize this latter, suppressive form of selective attention.

In our laboratory, we have been particularly interested in detailing the top-down mechanisms by which the brain establishes biased attentional states –that is, anticipatory sets. Until recently, the vast majority of attention research had concentrated on the effects of attention on the processing of attended versus unattended stimuli that occurs subsequently, as a result of attentional biasing (see Desimone and Duncan, 1995; Hillyard and Anllo-Vento, 1998). When the brain is involved in attentionally demanding tasks, it is advantageous to direct attention to the relevant stimulus, location or action that is to be engaged in, *prior to* the arrival of the stimulus or execution of the action. Our laboratory, among others, has reported a neural correlate of the suppressive form of biasing in a number of contexts, including suppression of one sensory modality in favor of another (Foxe et al., 1998; Fu et al., 2001), and suppression of portions of space that are likely to contain distractor information (Worden et al., 2000; Kelly, Lalor, Reilly and Foxe, 2006). The results of these studies motivated the current project, and the phenomenon of interest, increases in the 8-15 Hz “alpha” frequency band of the EEG, serves as one of our dependent measures of attentional biasing.

Using an instructional cueing paradigm, our lab originally found differential modulation of parieto-occipital alpha activity in the period preceding an imperative stimulus (S2) when visually presented word cues (S1) instructed subjects in an intersensory attention paradigm (Foxe et al., 1998). In that experiment, subjects were cued to attend either the visual or auditory portion of a subsequent compound audio-visual stimulus. The cue instruction to attend the auditory modality (the word “BEEP”) resulted in significantly larger 8-14 Hz activity over parieto-occipital cortex in the anticipatory period preceding the S2 than when subjects were cued to attend the visual modality (the word “FLASH”). This effect was not due to eye closing or gaze issues as eye movements were carefully monitored and subjects were instructed to maintain fixation during both attend-auditory and attend-visual anticipatory periods. It was hypothesized that this oscillatory enhancement in the case of the attend-auditory instruction might reflect anticipatory gating of visual processing by parieto-occipital structures, which are known to be involved in attentional switching and disengagement within the visual modality (e.g. Farah, Wong, Monheit and Morrow, 1989; Posner and Petersen, 1990; Posner, Walker, Friedrich and Rafal, 1984).

Foxe and colleagues (1998) interpreted their results in terms of visual spatial attention theories that posit parietal control of visual selective attention such as those of Van der Heijden (1991) or LaBerge (1997). In such models, the posterior parietal areas of the fast dorsal visual stream mediate the extent of processing of a particular stimulus through parietal inputs to the thalamic relay nuclei, nucleus reticularis and the putamen, and also through modulatory inputs to the ventral visual stream, which is specialized for the high level extraction of featural information from stimuli. That is, the spatially specialized dorsal stream areas control the gateways to ventral visual processing, according to theory. Our lab suggested that the enhanced parieto-occipital alpha-band activity seen in our intersensory selective attention paradigm represented a selective inhibition of "all" of visual space and was one mechanism by which visual space could be selectively inhibited. The prediction from these results would be that during visual spatial selective attention paradigms, such alpha-gating mechanisms should be selectively deployed to regions of space that need to be ignored. That is, alpha-gating should occur in cortices that are likely to receive *distracter* information. Our lab then examined this prediction in a visuo-spatial cueing study (Worden et al., 2000).

In that visuo-spatial experiment, centrally presented arrow cues were used to direct attention to the left or right visual fields and alternately to the upper and lower quadrants within the visual hemifields. Subjects were required to make orientation and motion-direction judgments upon the subsequent S2 stimuli. It should be emphasized that only stimuli occurring at the cued location were relevant to the task, such that stimuli occurring opposite to the cued location could be deemed distracters. Compellingly, deployment of visual attention to one hemifield resulted in anticipatory sustained focal increases in alpha-band activity over occipital cortex ipsilateral to the cued direction of attention. Further, these foci also moved retinotopically depending on whether attention was deployed to the upper or lower quadrants. Thus, enhanced alpha-band activity was seen over the retinotopic area that was likely to contain distracter stimuli, consistent with a role for these oscillatory processes in inhibition of non-cued spatial locations and consistent with the predictions of Foxe et al. (1998). This basic finding of retinotopically-specific alpha-band mediated suppression of potentially distracting locations has subsequently been replicated and expanded several times (Cosmelli et al., 2011; Kelly et al., 2006; Rihs, Michel and Thut, 2009; Sauseng et al., 2005; Thut, Nietzel, Brandt and Pascual-Leone, 2006; Yamagishi, Goda, Callan, Anderson and Kawato, 2005).

The research question that motivated the current project was whether this gating function of alpha-band activity was specific to spatial and intersensory attention, or rather if alpha could play a role in the biasing of attention between other stimulus parameters within a given modality, such as visual features (**Chapter 1**). It was well established from prior research that attention can be deployed to non-spatial visual features, such as color or motion parameters, facilitating the processing of subsequent stimuli incorporating the attended feature, independently of spatial location (e.g., Corbetta et al., 1991; Egner et al., 2008; Martinez-Trujillo and Treue, 2004; Most and Astur, 2007; Wylie, Javitt and Foxe, 2004, 2006). We again utilized an instructional cuing paradigm, although in this case the cues directed participants to attend to either the color or motion of an upcoming stimulus consisting of many small moving and colored dots. Participants were tasked with making a difficult discrimination with respect to the cued feature. We chose the features of color and motion because they have well-described, spatially-separated cortical processing centers in the dorsal and ventral visual streams, respectively (Ungerleider and Mishkin, 1982). We predicted that, if feature-based attention invokes alpha-band suppressive mechanisms analogously to spatial attention, then we should observe a dorsal-ventral shift in the source of alpha-band oscillations depending on which feature was attended. The results of this experiment are discussed in **Chapter 1**.

BINDING VERSUS SELECTIVE ATTENTION

As described above, the process of selective (or biased) attention aims to separate the environmental signals into discrete classes of information, such that the different classes of information can be enhanced or suppressed, depending on the current behavioral context. This process is largely endogenous, or “top-down.” Seemingly in opposition to this selective mode of operation, however, is an inherent tendency to integrate – or “bind” – stimulus features that co-occur in time and space. Previous literature has demonstrated that when one part of an object is attended, enhanced processing spreads from an object's attended features to its irrelevant features (Albrecht, List and Robertson, 2008; Egly, Driver and Rafal, 1994; Fiebelkorn, Foxe and Molholm, 2010a; Fiebelkorn, Foxe, Schwartz and Molholm, 2010b; Katzner, Busse and Treue, 2009; Molholm, Martinez, Shpaner and Foxe, 2007; O’Craven, Downing and Kanwisher, 1999; Schoenfeld et al., 2003; Wylie et al., 2004).

For example, Egly and colleagues (1994) showed participants a series of rectangle pairs. On each trial, one end of one of the rectangles would brighten momentarily to serve as a cue. The participants then needed to detect a target, which usually occurred at the cued location (valid), but would occasionally occur at the uncued end of the cued rectangle or at the equidistant end of the uncued rectangle (invalid). The key result was that participants were faster to respond to invalidly-cued targets occurring within the cued rectangle than to invalidly-cued targets occurring within the uncued rectangle. Since the distance from the cued location was identical in each case, the interpretation of this result is that attention is deployed across an entire object when just one of its parts is cued.

Besides automatically spreading across the spatial extent of an object, attention also spreads across the features of an object. For example, Schoenfeld and others (2003) directed attention to one or the other of two spatially overlaid stimuli. On some trials, one of these two stimuli underwent a task-irrelevant color change. Electrophysiological and neuroimaging results of that experiment indicated that the brain response to the color change was greater when the change occurred in the attended stimulus, notwithstanding that color was completely irrelevant to the task. Again, the interpretation is that the entire object receives enhanced processing, despite that only a subpart of the object is explicitly relevant.

This spreading of attention across the features of an object seems to be highly automatic (i.e., "bottom-up"). For example, a prior experiment by our group (Fiebelkorn et al., 2010a) found a spread of attention to task-irrelevant auditory stimuli simply because they co-occurred with an attended visual stimulus. This spread of attention occurred despite the fact that the sound was often semantically incongruent with the visual stimulus and often unrelated to the attended object category, which served to make the sound highly distracting from the visual-only task. Because attention to the sounds would disrupt task performance (due to incongruent pairings), the fact that attention persisted in spreading to the sounds suggests that this process is highly automatic.

When designing the experiment described in **Chapter 1,** we became concerned with the interaction between selective attention and binding processes. We reasoned that if attentional enhancements naturally spread from an object's attended feature to its task-irrelevant features, then such a spread of attention might mitigate the suppressive mechanisms of selective attention that we sought to measure. Thus, we tried to reduce the tendency to bind the features of the stimuli in that experiment by using briefly-presented random dot arrays in which each dot moved at a unique rate. In the experiments that followed, we revisited the prediction that binding processes would interfere with selective feature-based attention and tested it explicitly (**Chapters 2** and **3**).

In the experiments described in **Chapters 2** and **3**, we placed feature-based selective attention and bottom-up binding processes in direct conflict to test whether top-down biasing can supersede the object-based spreading of attention under conditions where binding of basic features is detrimental to task performance. To this end, we used feature-based cued attention tasks similar to the one used in **Chapter 1**. At the same time, we used well-established Gestalt grouping principles to manipulate the degree to which the imperative stimulus was perceived as a single coherent object (i.e., we differentially engaged binding processes). The optimal approach for our participants would have been to suppress the distracting (i.e., uncued) feature, which would require that feature-based selective attention disrupt or diminish bottom-up binding processes that typically occur in response to a coherent object. If binding is automatic, however, then we should rather see a diminution of selective attention measures when binding demands are increased.

For the experiment described in **Chapter 2**, we used a probabilistic cuing paradigm: participants were instructed to respond to all targets, including those that were invalidly-cued. This approach enabled us to use response time as a dependent measure, which is a classical and compelling measure of attentional biasing (e.g., Posner, 1980). However, the probabilistic paradigm does not call for suppressive attention, as all stimuli are potentially relevant. Because we hypothesized that the greatest degree of interference between selective attention and binding would occur when one feature of an object was explicitly *suppressed*, we repeated the task as an instructional paradigm (participants responded to cued targets only). In an instructional cuing paradigm, however, there is no overt measure of processing efficiency, and so for the experiment described in **Chapter 3**, we used event-related potentials (ERPs) to observe biased processing at the neural level. This approach moreover enabled us to characterize the timing and localization of the brain regions involved in resolving the conflict between top-down selective attention and bottom-up binding processes.

CHAPTER 1

Anticipatory Attentional Suppression of Visual Features Indexed by Oscillatory Alpha-Band Power Increases: a High-Density Electrical Mapping Study.

When covertly attending to regions of space where behaviorally relevant information is expected to occur, processing of visual stimuli appearing at those locations is enhanced (e.g., Hillyard, Vogel and Luck, 1998; McMains, Fehd, Emmanouil and Kastner, 2007). Conversely, if a region of space is expected to be a locus of distracting events, processing of stimuli occurring there is attenuated (e.g., Hillyard et al., 1998; Rees, Frith and Lavie, 1997). It is also clear that animals can use available information about the probable location of an upcoming relevant or distracting event to prepare their brains in advance, such that relevant information will receive enhanced processing whereas distracters will be suppressed (Foxe, Simpson, Ahlfors and Saron, 2005a; Luck et al., 1997; McMains et al., 2007). For visuospatial selective attention tasks, the suppressive aspect of such anticipatory preparation appears to be reflected in retinotopically specific transient increases of alpha-band (~8-15Hz) oscillatory power in the EEG (Kelly, Lalor, Reilly and Foxe, 2005, 2006; Rihs et al., 2007; Sauseng et al., 2005; Thut et al., 2006; Worden et al., 2000; Yamagishi et al., 2005).

Based on the cellular physiology of similar oscillations in animals, it has been proposed that alpha might serve as a functional gating mechanism (Foxe et al., 1998; Lopes da Silva, 1991). In a direct and compelling test of this gating role for alpha activity, Romei, Rihs, Broadbeck and Thut (2008a,b) stimulated visual cortex with transcranial magnetic stimulation (TMS) while monitoring alpha power. They found that the probabilty of subjects experiencing visual percepts (phosphenes) was inversely related to the amplitude of ongoing alpha activity in occipital cortex. That is, TMS was less effective at evoking visual percepts when alpha power was high, suggesting that the excitability state of these regions was relatively lower during higher alpha periods.

Furthermore, the network of brain areas that contributes to the generation of alpha rhythms, which includes frontal, parietal and occipital visual areas, as well as thalamic nuclei (Lindgren et al., 1999; Lopes da Silva, 1991), is implicated in several influential theories of attention (e.g. LaBerge, 1997; Posner and Petersen, 1990). The common theme in such models is that goal representations in frontal areas interact with parietal attentional control mechanisms to bias sensory processing, such that relevant information is preferentially processed while competing information is reduced. Providing strong evidence linking alpha-band oscillations to the fronto-parietal attention network, Capotosto, Babiloni, Romani and Corbetta (2009) showed that repetitive TMS to frontal or parietal sites disrupted the subsequent attentional modulation of alpha oscillations at occipital locations, and that such disruption was related to decrements in performance.

The critical role of alpha-band oscillations in selective attention has thus been clearly demonstrated. To date, however, the alpha-band effects of selective attention have only been characterized with respect to spatial and intersensory attention (e.g., Foxe et al., 1998; Fu et al., 2001; Gomez-Ramirez et al., 2007). Attention can also be deployed to non-spatial visual features, such as color or motion parameters, facilitating the processing of subsequent stimuli incorporating the attended feature, independently of spatial location (Corbetta, Miezin, Dobmeyer, Shulman and Petersen, 1991; Egner et al., 2008; Martinez-Trujillo and Treue, 2004; Most and Astur, 2007). Here we asked whether the role of alpha-band oscillations is specific to spatial and intersensory attentional selection, or if it is a more flexible system that also serves to suppress irrelevant features during feature-based selection. Our goal was to further characterize the alpha-band attentional measure by testing its spatiotemporal properties in a purely feature-based attention task. To the extent that alpha-band activity serves as a general attentional suppression mechanism, one would predict that alpha-band power shifts between feature-selective cortical regions processing irrelevant features analogously to the way in which alpha-band power shifts between retinotopic areas processing irrelevant parts of space. In designing this study, we chose to test selective attention between the features of motion and color specifically because processing of these features is localized to spatially disparate cortical regions in the dorsal and ventral visual streams, respectively (Ungerleider and Mishkin, 1982).

METHODS

Participants

Twelve adults (9 male, 2 left-handed) aged 21 to 50 years (mean: 30.5±8.2) participated in the experiment. Participants were sourced from the undergraduate and graduate student populations of The City College of New York, and from the local community. Eleven participants had normal color vision. One participant could not perform the color discrimination using the typical ‘long minus medium wavelength’ (L-M) axis of the Derrington, Krauskopf, and Lennie (DKL) color space (see section Stimuli, below). This participant performed the task using the ‘short minus long plus medium wavelength’ (S-[L+M]) axis of the DKL color space. Data were analyzed with and without this participant. Inclusion of this subject did not affect the pattern of group results, and the subject's overall pattern of results resembled that of the remainder of the group, and so this subject was included in the group analysis. None of the participants had any history of brain injury or disease, per self-report. All participants provided informed consent prior to the experiment. All materials and procedures were approved by the institutional review board of The City College of New York in accordance with the United States Public Health Service Act (US 45 CFR 46).

Cueing strategy

We employed a variant of the common S1-S2 cuing paradigm (e.g., Posner, 1980), in which a symbolic arrow cue (S1) directs attention to a part of space where subjects are to scrutinize a subsequent stimulus (S2) and indicate if it satisfies some target condition. In our case, features of the upcoming stimulus were cued, rather than spatial locations.

In the classical S1-S2 cuing task, the cues are probabilistic in nature. That is, the cue will likely indicate the correct location where the S2 will occur. However, subjects are instructed to respond to all targets, including those that occur at an uncued location. Non-informative (neutral) cues are also often included as a baseline condition. The typical finding is speeded responses to target stimuli following valid cues and slowed responses to targets following invalid cues, relative to neutral cues (e.g., Posner, 1980). This pattern of results is taken to indicate biasing of attention toward the cued location and away from the uncued location. However, when probabilistic cues are used and subjects are instructed to respond to all targets, there is no strategic impetus to suppress processing of uncued locations. Indeed, uncued locations are *relevant* and must be *attended* at least to some extent. The finding of a reaction time cost for invalidly cued targets is thus typically interpreted as *less enhancement* of processing relative to the neutral condition, as there is no *a priori* reason to suppose suppression of processing at that location, whereas some attentional enhancement would be advantageous. This is clearly not ideal for a study investigating a measure of attentional suppression (i.e., alpha). On the other hand, instructional cues (as in: Worden et al., 2000) direct subjects to respond only to targets occurring at the cued location and to ignore all events at uncued locations. In this case, potential stimuli appearing at uncued locations would in fact be *distracting*, and *suppression* of processing at those locations would be advantageous. However, in this case there is no behavioral metric of attentional processing, since the concept of 'cue validity' no longer applies. In pilot work, we employed a probabilistic cue to determine if attention can be selectively employed in our feature-based design as indicated by the standard reaction time measure, which was in fact the case (see also **Chapter 2**). We then used instructional cues for the EEG experiment investigating the alpha-band measure to encourage suppression of irrelevant features.

Stimuli

All stimuli were presented on a standard size cathode ray tube (CRT) monitor with a 75Hz refresh rate. Trials began with a warning cue consisting of a white fixation dot on a black background for 1s, followed by a cue word in white block capitals ('COLOR', 'HUE', 'MOTION', or 'DIRECTION') for 1s. The words 'color' and 'hue' both directed attention to the color of the stimulus. Likewise, 'motion' and 'direction' both directed attention to the motion of the stimulus. After an interval of 1.7 to 2.3s (random, and evenly distributed) during which only a black screen was displayed, the S2 was presented for 0.2s (**Figure 1**).

The S2 consisted of an array of one thousand dots, each subtending 0.05 degrees of visual angle, constrained to a square aperture subtending 5 degrees of visual angle. Each dot moved on a linear trajectory with a unique velocity of 18 to 36 degrees per second (evenly distributed). Dots ‘wrapped around’ the edges of the square aperture, so that the total amount of illumination was held constant.

Dots were typically colored with a hue from the L-M axis of an isoluminant plane of DKL color-space (Derrington, Krauskopf and Lennie, 1984), although one subject was unable to perform the task with these colors and instead used the S-(L+M) axis of the DKL color-space (see section Participants, above). This color-space uses the response properties of neurons in macaque lateral geniculate nucleus to create a subjective luminance axis, planes orthogonal to which are approximately isoluminant. The use of this color-space enables the continuous variation of hue needed to derive hue discrimination thresholds while controlling for subjective luminance.

Task

On standard trials, all dots moved in the same direction and had the same hue. On target trials, 20% of dots differed from the majority either by having a different trajectory or a different hue. No particular value of any feature indicated a target: subjects had to detect if any two values of the cued feature were present. This strategy was used to reduce competition within a feature processing area (if subjects were attending to red and suppressing green, for example). The goal, rather, was to have subjects attend to color and suppress motion, or vice versa as the cue indicated.

Targets and non-targets were equally likely (50%), and 17% of trials had targets in both features. In the case that a target was present in both features, the particular dots constituting the target for each feature were chosen independently. Subjects were instructed to respond with a button press as quickly as possible upon detecting a target if and only if that target occurred in the cued feature. Each S2 was followed by a 1s response interval. Each subsequent trial began immediately following the response interval.

Prior to beginning the experiment, performance was titrated to ~80% target detection rate for both direction discrimination and hue discrimination using an up-down transformed response (UDTR) modified staircase procedure (Wetherill & Levitt, 1965). After titration, subjects completed ten 10-minute blocks (with self-paced breaks after every 12th trial).

EEG Recording

Continuous EEG was acquired through the ActiveTwo BioSemi (Amsterdam) electrode system from 168 scalp electrodes, digitized at 512Hz. Active electrodes integrate the first amplification stage directly with the Ag/AgCl sensor, greatly reducing the effects of electronic noise. For practical purposes, the output impedance of the active sensor is less than 1Ω. With the BioSemi system, every electrode or combination of electrodes can be assigned as the reference, which is done purely in software after acquisition. BioSemi replaces the ground electrodes that are used in conventional systems with two separate electrodes: Common Mode Sense (CMS) active electrode and Driven Right Leg (DRL) passive electrode. These two electrodes form a feedback loop that drives the average voltage of the subject (i.e., the common mode voltage) as close as possible to the reference voltage of the analog-to-digital converter. Signals are recorded as the voltage between each electrode and the CMS. For a detailed description of the referencing and grounding conventions used by the BioSemi active electrode system, visit www.biosemi.com/faq/cms&drl.htm. During online data collection, signals were band-pass filtered between 0.1-100Hz. Data were re-referenced offline to the average activity and downsampled to 32Hz (see section Independent component analysis, below). EEG data were processed using the FieldTrip toolbox (Donders Institute for Brain, Cognition and Behaviour, Radboud University Nijmegen, the Netherlands. See http://www.ru.nl/neuroimaging/fieldtrip) for MATLAB (The MathWorks Inc., Natick, Massachusetts).

Independent component analysis (ICA)

Our goal was to examine the oscillatory activity within a delimited frequency band arising from different locations in the brain. A conventional approach to this question is to band-pass filter the scalp-recorded data at the frequency range of interest, and then estimate the inverse source solution of the filtered data. A more powerful approach is to first separate the data into components attributable to different sources, and filter subsequently. This latter approach has several advantages, most importantly: 1) source estimation after an operation such as spectral filtering has questionable validity, whereas filtering after source separation does not suffer this drawback (Graimann and

Pfurtscheller, 2006), and 2) it has been demonstrated that ICA can robustly separate artifacts from brain-related activity (e.g., Delorme, Sejnowski and Makeig, 2007).

We used the FastICA algorithm (Hyvärinen, 1999) to decompose each subject's data into independent components. We used a deflation approach to the fixed-point algorithm with a cubic nonlinearity and the following parameters: $\varepsilon = 10^{-4}$, $\mu = 1$, and 1000 iterations maximum. We did not use any additional algorithmic fine-tuning or stabilization. These were the default settings for the FastICA algorithm. The FastICA toolbox is available at http://www.cis.hut.fi/projects/ica/fastica. Subsequent analyses were performed on the isolated components. Because the raw datasets were extremely large ($3x10^6$ time-points x 168 channels), we first downsampled the timeseries to 32Hz sampling rate for computational tractability. This preserves frequency information below 16Hz, and thus does not impact the planned analysis in the 8-15Hz alpha frequency band.

Temporal spectral evolution (TSE) analysis

To examine the spatiotemporal dynamics of alpha-band amplitude in the cue-target interval, TSE waveforms were derived by the following method. First, epochs time-locked to the cue (200ms pre- to 3500ms post-cue) were band-pass filtered (third-order digital Butterworth zero-phase, 8-15Hz, 24dB octave). Second, the complex analytic signal for each component was derived by the Hilbert transform. Third, the instantaneous amplitude envelopes of the analytic signals were computed by taking the absolute magnitude of the complex waveforms. Fourth, the amplitude envelopes were baseline-corrected and averaged across trials. Over 250 sweeps were available for each average.

Identification of alpha-reactive components

We next identified the components that showed a change in alpha-band power in the cue-target interval for each subject. To do this, we took the average amplitude for each component in the last second of the cue-target interval (1.7-2.7s post-cue onset). We chose this interval because it begins late enough that the evoked response from the offset of the cue will have dissipated, but ends before contamination by the sensory response to the S2 has begun for any trial. The alpha-reactive components were defined as those that had an amplitude three standard deviations above (positively-reactive) or below (negatively-reactive) the mean of the set of all components in this window. We separately considered positively- and negatively-reactive components because these could have different functional interpretations. Alpha-reactive components were then tested for sensitivity to cuing condition.

Identification of feature-sensitivity of alpha-reactive components

For each alpha-reactive component of each subject, we performed a running two-sample t-test comparing the amplitude following a motion cue to the amplitude following a color cue. The criterion for significance was 30 consecutive time points with $p < 0.05$. The directionality of the effect was determined by summing the t-scores in the last second of the cue-target interval (1.7-2.7s post-cue onset). Negative values indicate that alpha power for the component was significantly greater for the attend-motion condition than the attend-color condition (motion-sensitive), and positive values indicate the converse relation (color-sensitive). Each alpha-reactive component with feature-sensitivity was then source-localized.

Source localization

To determine the localization of feature-sensitive alpha-reactive components, we performed source modeling using brain electric source analysis (BESA 5.1.8, MEGIS Software GmbH, Gräfelfing, Germany; Scherg and Von Cramon, 1985). BESA employs a least-squares fitting algorithm, defining location and orientation of dipoles for which the maximal amount of variance is explained (see Scherg and Picton, 1991; Simpson et al., 1995). For the purpose of modeling, an idealized four-shell ellipsoidal head model with a radius of 90 mm and scalp and skull thickness of, respectively, 6 and 7 mm was assumed. In most cases, the data were best explained by a pair of dipoles, one in each hemisphere. In two cases, the model was moderately improved by the addition of a third dipole to the approximate center of the shell. Since our hypotheses concern activation of visual cortices, only sources within this broad region of interest were retained. If a reasonable (>70% variance accounted for) model could not be attained with at most three dipoles, the component was rejected as physiologically implausible. This occurred for three components. In these cases, the respective subjects had other feature-sensitive components with larger effect sizes.

Statistical testing of source distributions

In line with our central hypothesis, we tested whether the distributions of sources for color-sensitive and motion-sensitive components had significantly different spatial means. That is, we wished to assess whether there was an obvious dorsal versus ventral stream bias to the spatial locations of these components. To perform this statistical comparison, we utilized a non-parametric bootstrapping procedure. This approach has the advantage that it makes no underlying assumptions about the population parameters of the distributions to be tested (e.g., normal distribution, homogeneity of variance, etc.). For the bootstrapping procedure, we first recorded the observed Euclidean difference between the centers of the two dipole groups within each hemisphere. Then we randomly repartitioned the dipole locations into two new groups and the Euclidean distance between the mean location of these groups was recorded for each hemisphere. This repartitioning procedure was iterated 10^4 times to create a distribution of inter-group distances that reflects random sampling. The statistical probability that the observed group difference is due to chance (i.e., the *p*-value) is the proportion of the distances from the bootstrapped distribution with a greater value than the observed distance.

Alpha-band power and behavioral performance

Presumably, the goal of preparatory attentional processes is to improve behavioral performance. Thus, alpha-band power increases, which are hypothesized to reflect attentional mechanisms, should bear some relation to behavioral performance. Specifically, if alpha-band power increases reflect suppression of potentially distracting information, then if alpha-power is not increased, the subject should be more distractible and therefore more prone to missing targets (see e.g. Kelly, Gomez-Ramirez and Foxe, 2009).

We separately analyzed components that showed greater alpha power for attention to color and those that showed greater alpha power for attention to motion. We also separately considered trials for which color was cued and those for which motion was cued. For each of these four combinations, we performed a median split of the single trials based on the average alpha power in the 200ms immediately preceding the S2, yielding “high alpha” and “low alpha” trials. We then compared the hit-rates for “high alpha” and “low alpha” trials for each of the four combinations of cue and feature-sensitivity.

We also tested the relationship of prestimulus alpha-band power to reaction time. For this analysis, we considered only correct positive responses ("hits"). For each component, we computed the Pearson product moment correlation coefficient between prestimulus alpha-band power (200 to 0ms pre-S2) and reaction time for each condition, attend-color and attend-motion.

RESULTS

Alpha-reactive components

ICA decomposition yielded between 152 and 167 independent components for each subject. Each subject had at least one positively-reactive component and two subjects each had one negatively-reactive component (**Table 1**). Examples of component TSEs for representative subjects are shown in **Figure 2**.

Feature-sensitivity of alpha-reactive components

Eleven of the twelve subjects had at least one positively-reactive component that had significant differences due to which feature was cued (**Table 1**). Examples of components showing such differences are shown in **Figure 3**. The remaining subject had one negatively-reactive component that had a significant difference due to which feature was cued. Because the single negatively-reactive component was unique, it is treated as a special case when presenting the results below.

One subject had one component that had greater amplitude when color was cued and two separate components that had greater amplitude when motion was cued. Seven of the remaining subjects had only components with greater power when color was cued, and four subjects had only components with greater power when motion was cued. We found that this dissociation was related to the subjects' discrimination thresholds for each feature type (see section Behavior, below).

Source localization

Dipole-equivalent estimations for component sources are plotted in **Figure 4** and summarized in **Table 2**. Components with greater power when color was cued were localized generally to dorsal visual stream regions, whereas components with greater power when motion was cued were localized generally to ventral visual stream regions. While it is sometimes considered that deep and ventral sources are difficult to detect with EEG, it has been demonstrated that these sources can be readily observed if ICA is first applied to isolate the source topography (e.g., Onton and Makeig, 2006). The spatial distribution of sources we observed is consistent with alpha-band increases reflecting suppression of processing of the to-be-ignored feature. The two distributions of sources had significantly different centers in both the left ($p = 0.0070$) and right ($p = 0.0028$) hemispheres.

The negatively-reactive feature-sensitive component had greater alpha power (i.e., less decrease) when color was cued compared to when motion was cued. This component was localized to left parietal cortex (approximate Talairach coordinates: x=-34.8, y=-69.9, z=35.9). The variance accounted for (VAF) by this model was 59%. The addition of a second dipole to the model also localized to left parietal cortex and did not substantially improve the VAF (62%).

Behavior

We found no differences in hit-rate between “high alpha” and “low alpha” trials for any of the combinations of cue and feature-sensitivity (all *p*’s > 0.71). This result could simply be type II error or, alternatively, it could be that subjects are consistently correctly preparing mechanisms indexed by alpha-band measures and performance decrements are reflected by some mechanism not assessed by this study (e.g., changes in another frequency band).

The majority of components did not show a relationship between prestimulus alpha power and reaction time for either attention to color or attention to motion. However, five components showed small positive correlations (r's = 0.14 to 0.24. all *p*'s < 0.047). Of these, four components were color-sensitive (i.e., they generally had higher power when color was cued), and one was motion-sensitive. All correlations were positive regardless of which feature was cued. This suggests that these effects may be due to general arousal rather than factors of feature-sensitivity.

Nevertheless, we did find a relationship between discrimination thresholds and feature-sensitivity of alpha-reactive components. As mentioned above, we observed that most subjects had alpha-reactive components that increased amplitude only in response to one type of cue, either color or motion. We asked if this could be related to their ability to make each type of discrimination. While performance was pre-titrated to 80% for each feature type for each participant, the degree of difference between values for each feature needed to achieve this rate of performance varied across subjects. We found that those subjects having only components that increased alpha power when color was cued tended to have lower motion thresholds than color thresholds, in terms of percent-of-maximum-possible (PMP) difference. Conversely, those subjects having only components that increased alpha power when motion was cued tended to have lower color thresholds than motion thresholds (**Figure 5**).

The single negatively-reactive component was localized to dorsal regions. This component had a greater decrease in alpha power when motion was cued compared to when color was cued, consistent with a suppressive role for alpha activity, and analogous to selective motion suppression. The subject with this component had larger color thresholds than motion thresholds (10.8 vs. 8.1 PMP, respectively), which is consistent with the overall pattern from the remaining subjects.

DISCUSSION

We found that when most people selectively deployed anticipatory attention to one feature of an upcoming stimulus array, alpha-band components of their EEG sharply increased in amplitude in the preparatory period. A subset of these components increased differentially depending on which feature was relevant. Components with greater alpha power when color was relevant (and motion was irrelevant), localized to more dorsal visual stream regions. In contrast, components with greater alpha power when motion was relevant (color was irrelevant), localized to ventral visual stream regions. Insofar as motion-processing is generally a more dorsal visual stream process and color-processing is generally a more ventral visual stream process, this pattern of results supports our thesis that such alpha-band increases reflect suppression of processing of the to-be-ignored feature. This is consistent with the hypothesized role of alpha-band increases in spatial attention, which have consistently been shown to index suppression of to-be-ignored parts of the visual field. Thus, alpha-band increases as a measure of attentional suppression are not specific to spatial attention, but also appear to operate during purely feature-based selection.

One participant showed an idiosyncratic result with an alpha-component that decreased in a feature-specific fashion. Specifically, the decrease was greater when motion was cued. This pattern suggests tonic suppression of motion-processing throughout the experiment, with phasic disengagement of suppression when motion became relevant. In line with this interpretation, this component was localized to a dorsal parietal region. While this pattern of results suggests a different cognitive strategy for this participant, it is still consistent with a suppressive role for alpha activity.

Since alpha-band increases appear to reflect suppression processes, those subjects that only had components that increased alpha power when color was cued (and motion was irrelevant) could be considered selective "motion suppressors" (note that the single subject with a negatively reactive feature-sensitive component fits naturally in this group). Conversely, those subjects that only had components that increased alpha power when motion was cued (and color was irrelevant) could be considered selective "color suppressors". We found that motion suppressors had lower thresholds for motion discrimination than for color discrimination whereas color suppressors had lower thresholds for color discrimination than motion discrimination. In other words, subjects appear to have selectively suppressed the "easier" feature when attending to the "harder" feature. One plausible interpretation of this result might be that when a given feature is particularly effortful to discriminate, then differences in that feature are unlikely to "pop out" and cause distraction to the subject, and therefore might not need additional active suppression. Another interesting perspective that has been suggested is that the oscillatory “architecture” of an individual’s brain leads to idiosyncrasies of performance (e.g., Hanslmayr et al., 2007; Romei et al., 2008a). In other words, it is the differential ability of the subject to engage suppression mechanisms that sets performance thresholds, rather than vice versa. The two interpretations are not mutually exclusive since it seems a reasonable proposition that a symbiotic development of perceptual and attentional processes could drive this relationship.

Other researchers have also examined this issue of oscillatory activity as it relates to feature-based selection. For example, Zanto and Gazzaley (2009) examined the EEG during the maintenance interval of a delayed-match-to-sample (DMS) working-memory task. For this task, subjects were presented with random dot stimuli similar to those in the present study and instructed to remember either the color, motion direction, or both. After an interval, a probe stimulus was presented and subjects responded if the probe matched one of the sample stimuli along the relevant feature dimension. In addition to stimulus-evoked broad-band potential measures, Zanto and Gazzaley also examined induced oscillatory activity in the maintenance interval in the alpha, beta and gamma frequency bands. Their main finding was that beta-band coherence was related to working memory performance, but they also observed clear alpha-band power increases over midline parietal scalp towards the end of the maintenance period, presumably reflecting preparatory activity. However, they did not find differences in alpha power based on which feature was relevant, nor did they find a connection to working memory performance. However, these investigators did not assess potential topographic differences in alpha, and task-performance levels were at or near ceiling such that attentional load and the need for suppressive processing were likely minimal.

Similarly, using magnetoencephalography (MEG), Jokisch and Jensen (2007) examined delay-period alpha-band activity during working memory maintenance, where subjects were required to recall either the identity or orientation of a face. Alpha-band power was found to be greater in dorsal regions when face identity (putative ventral stream information) was relevant, than when face orientation (putative dorsal stream information) was relevant, consistent with the results presented here. However, unlike the current study, alpha-band power was not found to be greater in more ventral areas when dorsal visual stream information (i.e., orientation) was relevant. In fact, no reliable sources of alpha were found for the orientation condition. There are several differences between the current study and that of Jokisch and Jensen that could account for this latter difference in findings. Firstly, of course, there are substantial differences in the nature of the tasks employed. Our study used an S1-S2 cuing paradigm whereas theirs used a DMS working memory task. Presumably, many processes characterize the maintenance interval of a DMS task, including encoding, maintenance, and preparatory processes. While the S1-S2 cuing paradigm has some working-memory component, subjects are not required to encode and maintain the same features they are later asked to evaluate in the S2. That is, subjects only need to maintain the instructional value of the cue, and not its color, for example. Thus, the S1-S2 cuing paradigm could be considered a purer assessment of feature-based preparatory processes. Further, in the DMS task, particular feature values are relevant, leading to potential competition within a given functional area processing the relevant feature dimension. For example, if the task is to attend for the color red and to ignore green, then both enhancement and suppression processes could both be invoked within the same color-processing region. Another aspect of the Jokisch and Jensen study seems germane here too. In their study, there were significant differences in performance between the two tasks such that the identification task was more demanding than the orientation task, with subjects significantly more accurate and faster for the latter. Recall that the alpha increases were seen over the dorsal stream during performance of the identification task – that is, there was increased suppression of the easier orientation task. Thus, as in our study, it was the easier task-feature, the one more likely to “pop out” as a distracter that was specifically suppressed. In turn, since subjects showed greater than 90% performance accuracy for the easier orientation task, perhaps strong suppression of the identification task was not invoked.

We did not find a reliable relationship between alpha-band power and performance accuracy on a single-trial basis. However, many studies have shown that such relationships between alpha-band power and performance accuracy exist for spatial attention (Wyart and Tallon-Baudry, 2009; Kelly et al., 2009). Furthermore, an inverse relationship between occipital alpha-band power and visual awareness of a near-threshold stimulus has been demonstrated outside of spatial attention tasks (Romei et al., 2008a; van Dijk, Schoffelen, Oostenveld and Jensen, 2008; Mathewson, Gratton, Fabiani, Beck and Ro, 2009). We did however find a small proportion of components that had a positive correlation between prestimulus alpha-band power and reaction time. However, this relationship was not dependent on which feature was cued or the feature-sensitivity of the component, suggesting that such a correlation reflects more general effects of arousal rather than feature-based selective attention. Furthermore, most components did not show any relationship between prestimulus alpha-band power and reaction time. Nevertheless, such a relationship has been demonstrated in a spatial selective attention task by Capotosto and colleagues (2009).

That we did not find a relationship between feature-sensitivity of prestimulus alpha-band power and speed or accuracy of performance on a single-trial basis may reflect essential differences between spatial and feature-based attention. One major difference between these forms of attention is that feature-based attention is characterized by gain modulation and sensitivity tuning whereas spatial attention is characterized by gain modulation alone (Ling, Liu and Carrasco, 2009). As direct physiological evidence of this, Martinez-Trujillo and Treue (2004), recorded from neurons in monkey MT while the monkey attended to the direction of motion of random dot stimuli. They found enhanced responses to stimuli moving at or near the attended direction of motion, but suppressed responses to stimuli moving in directions that differed greatly from the attended direction. If such suppression within an area processing the attended feature is mediated by alpha-band mechanisms, then this could potentially obscure the relationship of prestimulus alpha-band power to subsequent performance. We sought to minimize the effects of suppression within the areas processing the to-be-attended feature by making all values of the to-be-attended feature relevant (i.e., all directions of motion were relevant because the subject had to detect any two directions of motion). However, some residual suppression within areas processing the to-be-attended feature could remain. Indeed, while we observed that components showed greater increases for attention to a particular feature in comparison to attention to the other feature, alpha-band power increases relative to baseline occured regardless of which feature was cued. This observation could be related to the suppression observed by Martinez-Trujillo and Treue at the single-cell level.

A great deal of research has shown that there is a strong bias to attend to objects as wholes (e.g. Triesman, 2004; Blaser et al., 2000; Martinez et al., 2006; Molholm et al., 2007). That is, it is clear from many studies that when attention is directed to one feature of an object (e.g., its direction of motion), other constituent features of that object (e.g., its color) are also preferentially processed, even when those features are completely irrelevant to the task at hand (O'Craven et al., 1999; Schoenfeld et al., 2003, 2007; Wylie et al., 2004), presumably as a result of feature binding. To study a purely feature-based attentional mechanism, this bias to bind must be overcome. Because we wished to study pure feature-based attention, we discouraged the subjects from attaching the features to any object or location by having the stimulus display fill the screen and having each dot in the display move at an idiosyncratic speed and by requiring the subject to make a discrimination that cannot be performed by attending to a single dot. In this regard, our stimulus design allowed for the two task-relevant features to be treated independently by virtue of the fact that they were not naturally related to each other within an obviously identifiable object. In contrast, many previous studies of feature-based attention have used coherent motion dot arrays with uniform color such that the motion and color features tend to cohere as a moving transparent surface (e.g. Liu, Slotnick, Serences and Yantis, 2003; McMains et al., 2007; Stoppel at al., 2007). We believe that configuring the stimuli so that subjects can orient to individual features with minimal object binding may have been a critical factor in our observation of suppression processes. This aspect of our design may well be why we have been able to observe bidirectional alpha-suppression effects.

Table 1

Summary of components. For each subject, the number of components for which alpha power increased (or decreased) by at least 3 standard deviations in the last second of the cue-target interval are shown in the second column (components that decreased are given in parentheses). The last two columns contain the number of components with feature-based effects.

Subject	Alpha-reactive components	Feature-sensitive components	
		color>motion	motion>color
1	7 (1)	1	0
2	4	0	3
3	3	0	1
4	3	1	0
5	2	1	0
6	5	1	0
7	4	1	0
8	6	1	2
9	3	1	0
10	2	0	1
11	1(1)	(1)	0
12	2	0	2

Table 2

Summary of component source localization. Positions are given in approximate Talairach coordinates. Positively-reactive components only are listed. VAF: variance accounted-for. †: VAF reported includes a third dipole.

	Left Hemisphere				Right Hemisphere				Total VAF%
	X	Y	Z	VAF%	X	Y	Z	VAF%	
Color > Motion	-19.5	-93.0	14.9	65.4	19.5	-93.0	14.9	81.6	95.8
	-41.6	-33.1	58.0	41.1	0.5	-81.4	64.1	55.2	86.9
	-33.0	-63.9	30.1	61.6	33.0	-63.9	30.1	59.1	87.6
	-35.9	-39.3	31.2	0.6	35.9	-39.3	31.2	84.4	96.6
	-39.9	-75.5	-1.4	34.6	34.2	-70.6	7.0	71.9	99.1
	-16.9	-92.0	22.8	52.9	16.9	-92.0	22.8	81.3	94.0
	-26.6	-58.7	35.5	16.2	26.6	-58.7	35.5	84.6	96.9
Motion > Color	-61.9	-72.3	-19.9	22.6	61.9	-72.3	-19.9	40.0	77.8†
	-28.7	-86.0	6.2	52.7	34.3	-81.4	15.5	32.4	92.9
	-33.6	-85.9	0.0	23.8	37.3	-71.3	0.1	56.8	92.6
	-41.8	-41.1	9.2	49.1	47.1	-61.0	9.5	40.0	93.2†
	-33.2	-66.5	9.2	82.0	33.2	-66.5	9.2	65.2	90.9
	-46.7	-63.7	-1.8	68.8	31.8	-52.0	17.0	69.3	89.6
	-21.8	-65.0	-24.0	73.8	21.8	-65.0	-24.0	46.7	91.9
	-29.0	-86.7	8.7	32.1	29.0	-86.7	8.7	68.6	74.0
	-36.5	-66.5	4.5	51.1	36.5	-66.5	4.5	18.0	83.2

Figure 1 (caption): Schematic of task procedures. For each trial, subjects first viewed a fixation dot for 1s, followed by a cue word in block capitals for 1s. The cue (S1) was followed by an interval of 1.7-2.3s with no stimulation. After the cue-target interval, the random dot stimulus (S2) was shown for 0.2s, followed by a 1s response period. The next trial began immediately following the response interval. The arrow in the S2 represents the motion of the dots and was not actually present in the stimulus. Details have been enhanced for clarity of illustration. Timeline not to scale.

Figure 1:

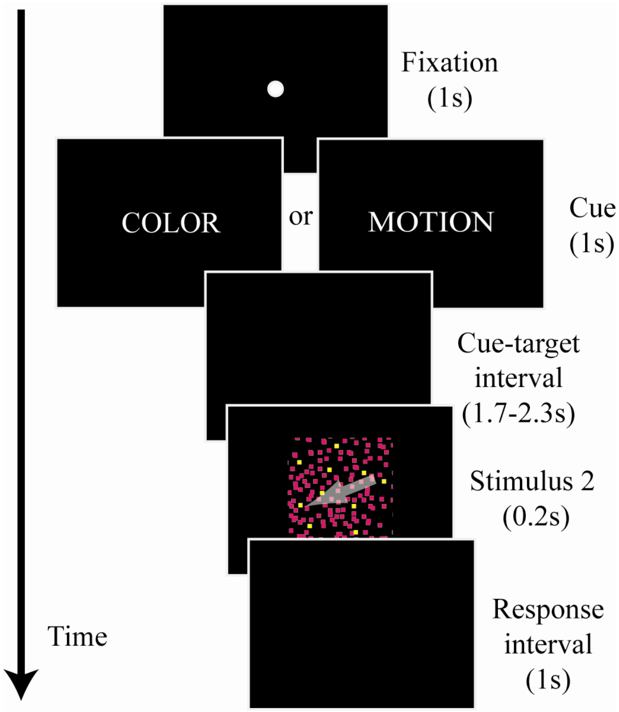
Fixation
(1s)
COLOR
or
MOTION
Cue
(1s)
Cue-target
interval
(1.7-2.3s)
Stimulus 2
(0.2s)
Response
interval
(1s)
Time

Figure 2 (caption): Temporal spectral evolution (TSE) waveforms for independent components. Sample data are shown for two subjects in the cue-target interval. Each trace is a TSE waveform showing the average time-course of induced amplitude in the 8-15Hz alpha band for one independent component. Arrows indicate ‘alpha-reactive’ components, which increase by more than three standard deviations in the final second of the preparatory period.

Figure 2:

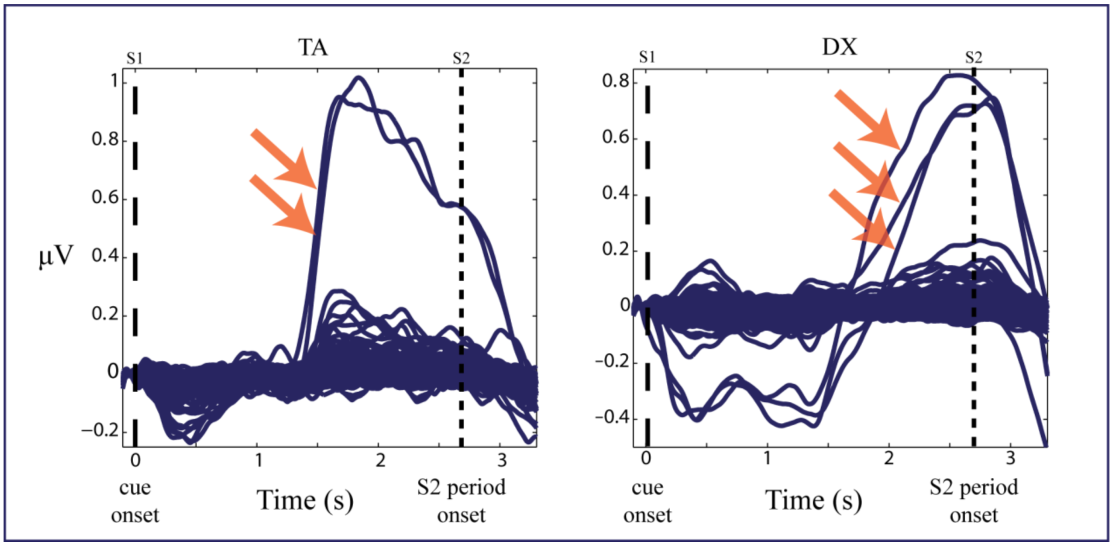
TA
S1
S2
μV
1
0.8
0.6
0.4
0.2
0
−0.2
0
1
2
3
cue onset
Time (s)
S2 period onset
DX
S1
S2
0.8
0.6
0.4
0.2
0
−0.2
−0.4
0
1
2
3
cue onset
Time (s)
S2 period onset

Figure 3 (caption): Feature-sensitive alpha-reactive components. Sample data are shown for two subjects in the cue-target interval. Left: TSE waveforms for attention to color (dashed) and motion (solid), with corresponding *p*-values for the difference between the two conditions shown below in orange. For the *p*-value plot the dashed line indicates a value of $p = 0.05$. Right: point-equivalent dipole source localizations for the given components.

Figure 3:

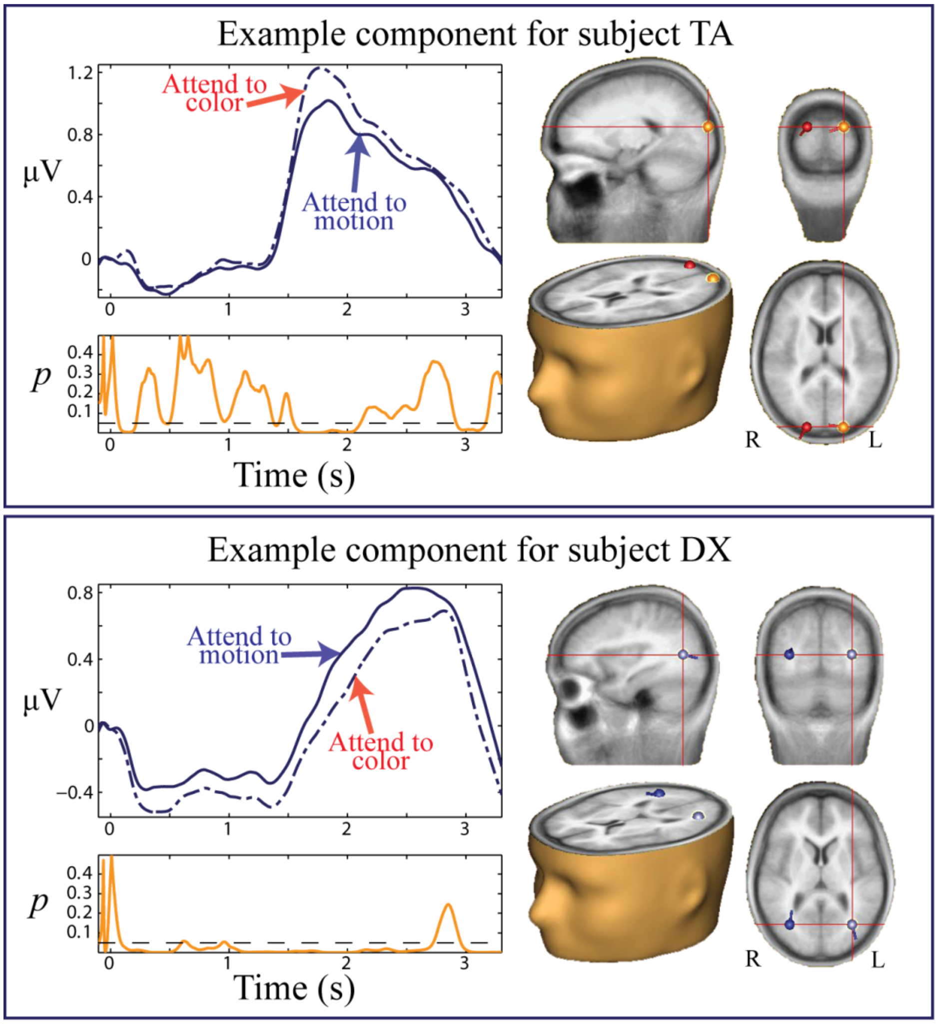
Example component for subject TA
1.2
0.8
0.4
0
μV
Attend to color
Attend to motion
0
1
2
3
p
0.4
0.3
0.2
0.1
Time (s)
R
L
Example component for subject DX
0.8
0.4
0
−0.4
μV
Attend to motion
Attend to color
0
1
2
3
p
0.4
0.3
0.2
0.1
Time (s)
R
L

Figure 4 (caption): Source localizations for all feature-sensitive alpha-reactive components. Components with greater alpha power for attention to motion than for attention to color are plotted in blue and are generally localized to ventral visual stream regions. Components with greater alpha power for attention to color than for attention to motion are plotted in red and are generally localized to dorsal visual stream regions. Note that the number of dipole pairs exceeds the number of subjects because a subject may have had more than one component that was feature-sensitive.

Figure 4:

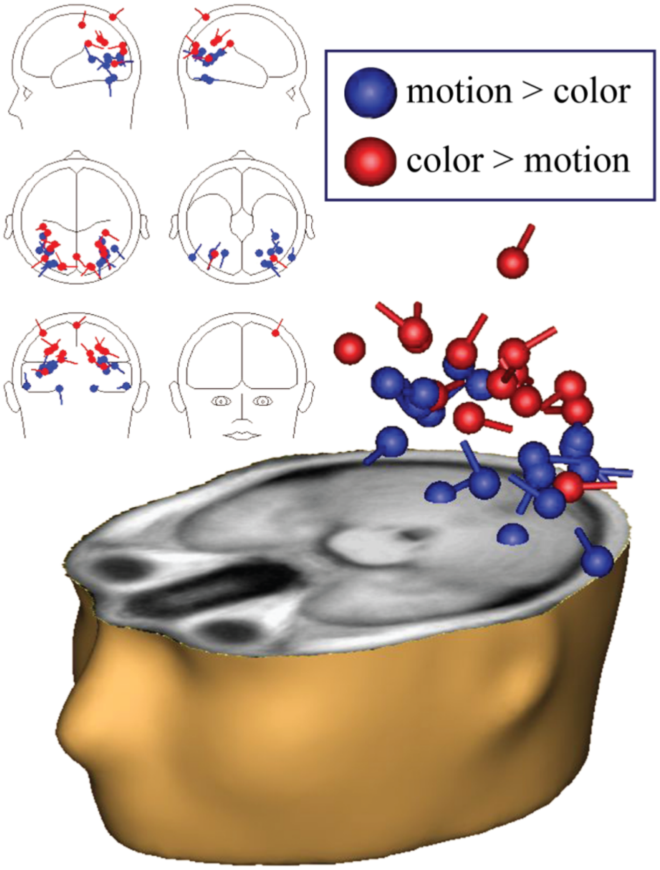

Figure 5 (caption): Motion direction and hue discrimination thresholds for motion and color suppressors. Units are the percent of the maximum possible difference between feature values. Subjects having only components that increased alpha power when motion was attended relative to when color was attended (color suppressors) tended to have higher discrimination thresholds for motion than for color. Conversely, subjects having only components that increased power when color was attended relative to when motion was attended (motion suppressors) tended to have higher discrimination thresholds for color than for motion. Error bars indicate standard error of the mean.

Figure 5.

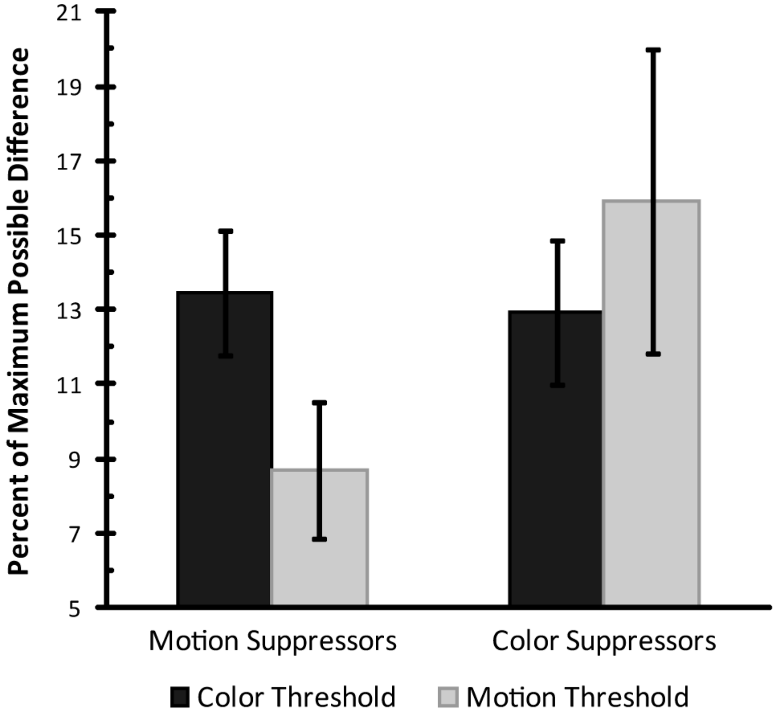
Percent of Maximum Possible Difference
21
19
17
15
13
11
9
7
5
Motion Suppressors
Color Suppressors
■ Color Threshold
□ Motion Threshold

CHAPTER 2

The Countervailing Forces of Binding and Selection in Vision.

Individuals with focal cortical lesions frequently present with highly-specific impairments of processing for one type of stimulus feature while other perceptual processes remain largely unaffected (e.g., Tamler, 1970; Vaina, 1994). These observations highlight a fundamental property of the brain's operation: the various forms of information available in complex environmental signals are processed in parallel channels that rely on distinct neural circuitry (see also Desimone and Ungerleider, 1986; Livingstone and Hubel, 1988; Tanaka, 1996). Information from these separate processing streams is then *synthesized* to engender a coherent experience characterized by the perception of discrete objects defined by conjunctions of many features. This process of binding together features also has some anatomical specificity, as evidenced by the striking difficulties of some individuals with parietal lesions. These patients have a spared ability to detect the presence of single features, but show impaired ability to identify particular conjunctions of features (Eglin, Robertson and Knight, 1989, 1991; Eglin, Robertson, Knight and Brugger, 1994; Esterman, McGlinchey-Berroth and Milberg, 2000; Laeng, Brennen and Espeseth, 2002; Pavlovskaya, Ring, Groswasser and Hochstein, 2002). Often this impaired binding is accompanied by impaired spatial processing and spatial attention, supporting the view that spatial attention is an important component of binding processes (cf. Treisman and Gelade, 1980). Indeed, stressing spatial attention demands increases the occurrence of illusory feature conjunctions in healthy research participants (Treisman and Schmidt, 1982).

In addition to their dependence on spatial attention, binding processes have been known to interact with feature-based attention. It has often been observed that when attention is directed to one feature of an object there is a tendency for task-irrelevant features of that object to receive enhanced processing also (Fiebelkorn et al., 2010a, 2010b; Katzner et al., 2009; O'Craven et al., 1999; Molholm et al., 2007; Schoenfeld et al., 2003; Wylie et al., 2004). The evidence therefore supports a hierarchical model of these processes whereby spatial attention precedes binding, which itself leads to a spread of feature-based attention.

Recently, our group performed the first investigation of the role of alpha-band oscillations in visual feature-based selective attention (**Chapter 1**; Snyder & Foxe, 2010). That study aimed to expand on prior research in selective attention which had shown that alpha-band power increases index the degree of attentional suppression in visuospatial attention tasks (Gomez-Ramirez, Kelly, Montesi and Foxe, 2009; Kelly et al., 2006, 2009; Kelly, Foxe, Newman and Edelman, 2010; Rihs et al., 2007; Sauseng et al., 2005; Thut et al., 2006; Worden et al., 2000; Yamagishi et al., 2005), intersensory attention tasks (Foxe et al., 1998; Fu et al., 2001; Gomez-Ramirez, et al., 2007), sustained visual attention tasks (Dockree, Kelly, Foxe, Reilly and Robertson, 2007; O'Connell et al., 2009) and audiospatial attention tasks (Kerlin, Shahin and Miller, 2010; Banerjee, Snyder, Molholm and Foxe, 2011). We asked if alpha-band mediated suppression would be observed in a feature-based selective attention task. In designing that experiment, we became concerned with the interaction of selective attention and binding processes. We reasoned that if spread of feature-based attention was an essential component of object-perception, such a spread of attention might reduce the suppressive mechanisms we sought to measure. Thus, we aimed to reduce the tendency to bind the features in our imperative stimuli by using briefly-presented random dot arrays in which each dot moved at a unique rate; a design with no drawback for that study if the underlying assumption (essential cross-feature spread of attention) was false.

Here, we explicitly test this prior assumption that facilitatory attentional selection is essential for binding. To do this, we specifically designed stimuli intended to place attention and binding demands in direct conflict with each other. Whereas prior investigations of the interaction of spatial attention with binding processes have manipulated attention and observed the effects on binding, we took here the converse approach of manipulating binding demands and observing the effects on selective attention. This approach was in line with the causal role that binding processes seem to have with regards to the spread of feature-based attention within objects.

We used a behavioral paradigm with arrays of dots characterized by one of two types of motion. In one case, all dots moved at the same speed and in the same direction. This type of motion strongly invokes the Gestalt principle of “common fate” (Wertheimer, 1923), leading to a percept of the dots as painted on a moving transparent surface. We termed this condition the STRONG object condition. For the other class of stimuli, all dots moved in the same direction, but each dot had a unique speed. This condition we termed the WEAK object condition, since the percept of a single surface is destroyed when the dots move independently. Our central thesis was that STRONG stimuli would invoke greater binding processes for object perception, and that such binding would counteract suppressive preparatory attention, as evidenced by response time (RT) measures. While increased binding due to Gestalt grouping has not been explicitly demonstrated, we deemed this was a reasonable assumption due to the clearer object-nature of the STRONG stimulus.

This paradigm was a variant of the classic cued-attention task. In a cued-attention task, a symbolic cue (“stimulus one”: S1) at the beginning of each trial provides task-relevant information about the nature of an upcoming imperative stimulus (S2). While the cue information is often helpful, or valid, on occasion the cue is misleading, or invalid. Non-informative neutral cues are included as a basis for comparison. For this experiment, cues directed attention to either the color or motion of the S2, and participants had to make a discrimination regarding one of those features. A typical finding is that valid cues lead to faster RTs compared to neutral cues whereas invalid cues lead to slower RTs compared to neutral cues (Posner, 1980). This pattern of results is taken to indicate that participants are able to utilize the cue information to prepare attention to facilitate the processing of the information that the cue had indicated while suppressing the processing of competing information.

In this case, the status of the imperative stimulus as an object was irrelevant to the task. Thus, if attentional selection was non-essential for binding in the face of a STRONG stimulus, then the optimal strategy would be to maintain suppression of irrelevant features. However, if attentional selection was essential for binding, then the presence of a STRONG stimulus would reduce suppressive attention effects compared to those with respect to a WEAK stimulus. We used RT as a dependent measure to test these predictions.

Specifically regarding the hypothesis that attention was essential for binding, we predicted the following:

1) For WEAK stimuli, valid cues would result in speeded RTs compared to neutral cues (RT benefit) and invalid cues would result in delayed responses (RT cost). This would indicate that attentional resources were allocated in favor of processing the cued feature at the expense of processing the uncued feature, which is what we mean by “attentional biasing,” in line with the findings of Posner and colleagues (1980) for cued spatial attention.

2) Compared to WEAK stimuli, RT benefits for STRONG stimuli preceded by valid cues would be reduced. This is because the anticipatory attentional suppression of the irrelevant feature would have to be counteracted to support binding for object perception.

3) Compared to WEAK stimuli, RT costs for STRONG stimuli preceded by invalid cues would be unchanged. This is because the anticipatory state would be equally biased against processing the invalidly-cued target feature regardless of the stimulus’ status as an object (i.e. STRONG or WEAK).

METHODS

Participants

Sixteen adults (8 male, 3 left-handed) aged 20 to 47 years (mean ± s.d.: 27.58±7.27) participated in the experiment. Participants were sourced from the undergraduate and graduate student populations of The City College of New York, and from the local community. None of the participants had any history of brain injury or disease, per self-report. Participants had normal or corrected-to-normal visual acuity per self-report. Fifteen participants had normal color vision. One participant was unable to discriminate red and green hues, but was able to perform the task normally using blue/yellow discrimination. All participants provided informed consent prior to the experiment. All materials and procedures were approved by the institutional review board of The City College of New York in accordance with the United States Public Health Service Act (US 45 CFR 46).

Experimental Design

Each participant completed two experiments involving judgments of random dot kinematograms. In one experiment, each dot in the stimulus moved at a unique speed (WEAK condition). In the other experiment, all dots in the stimulus moved at the same speed (STRONG condition). Experiments were administered in a counter-balanced order across participants. Participants completed the experiments in two separate sessions divided by an interval of at least two weeks. The purpose of this separation was to reduce task-set interference effects. Due to attrition, four participants completed only the WEAK condition and three participants completed only the STRONG condition. The remaining nine participants completed both conditions.

Task

Procedures were adapted from a prior feature-based attention experiment (**Chapter 1**; Snyder and Foxe, 2010). We used a cued-attention paradigm in which visual word forms (‘COLOR’, ‘HUE’, ‘MOTION’, or ‘DIRECTION’) indicated to the participant the relevant feature to attend. A non-informative neutral cue (‘NOTHING’) was used for comparison. The words ‘color’ and ‘hue’ both directed attention to the color of the stimulus. Likewise, ‘motion’ and ‘direction’ both directed attention to the motion of the stimulus. The use of multiple cue words for each feature was to reduce the automatization of the task through implicit learning, with an aim to maintain the engagement of endogenous orienting mechanisms throughout the session. Participants were instructed to use the cue information to guide them in detecting a target in a subsequent imperative stimulus (S2), and upon detection of a target to respond with a button press as quickly as possible without sacrificing accuracy. Targets were present on 50% of trials, and were characterized by a particular property in either feature dimension –described below, and illustrated in **Figure 6**. Twenty percent of targets were preceded by a neutral cue. Of informatively-cued targets, 80% were preceded by a valid cue (the cued feature was the feature that defined the target), and the remaining 20% were preceded by an invalid cue (the target was defined by the uncued feature). Participants were instructed to respond to all targets, including those that were invalidly cued.

Stimuli

The experiment was administered in a light- and sound-attenuated chamber using Presentation software version 14.4 (Neurobehavioral Systems). All stimuli were presented on a standard size cathode ray tube (CRT) monitor with a 60Hz refresh rate. Trials began with a warning cue consisting of a white fixation dot on a black background for 1s, followed by a cue word in white block capitals for 1s. After an interval of 1.7 to 2.3s (random, and evenly distributed) during which only a black screen was displayed, the S2 was presented for 0.2s. Each S2 was followed by a 1s response interval. Each subsequent trial began immediately following the response interval (**Figure 1**).

The S2 consisted of an array of one thousand dots, each subtending 0.05 degrees of visual angle, constrained to a square aperture subtending 5 degrees of visual angle. In the WEAK condition, each dot moved at a unique speed between 14 and 28 degrees of visual angle per second (evenly distributed). In the STRONG condition, all dots moved at a speed of 21 degrees of visual angle per second. Dots ‘wrapped around’ the edges of the square aperture, so that the total amount of illumination was held constant.

Dots were colored with hues from an isoluminant plane of DKL color-space (Derrington et al., 1984). This color-space uses the response properties of neurons in macaque lateral geniculate nucleus to create a subjective luminance axis, planes orthogonal to which are approximately isoluminant. The use of this color-space enables the continuous variation of hue needed to derive hue discrimination thresholds while controlling for subjective luminance.

Targets

For both WEAK and STRONG conditions, standard trials were characterized by all dots moving on a common linear trajectory and having the same color. For color targets in both the WEAK and STRONG conditions, 20% of dots had a different color than the majority while the dots continued to move as on standard trials for the respective condition (i.e., WEAK or STRONG). For motion targets in the WEAK condition, 20% of the dots moved on a different linear trajectory than the majority. For motion targets in the STRONG condition, all dots moved on a common curved trajectory. Schematized examples of targets are illustrated in **Figure 6**. The degree of difference for each of the relevant targets was titrated on a per-subject basis to 80% detection rate prior to beginning each experiment using an up-down transformed response (UDTR) modified staircase procedure (Wetherill and Levitt, 1965). No particular value of any feature indicated a target: subjects had to detect a particular feature *variation* in the stimulus. This strategy was used to reduce competition within a feature processing area (if subjects were attending to red and suppressing green, for example). The goal, rather, was to have subjects attend to *color* and suppress *motion*, or vice versa as the cue indicated.

Analysis

Response time (RT) data were analyzed using SPSS Statistics software version 17.0. We used a 3x2x2 univariate analysis of variance (ANOVA) with random factors of ‘validity’ (three levels: invalid, neutral and valid), ‘feature’ (two levels: color and motion) and ‘object’ (two levels: STRONG and WEAK). The factor of ‘subject’ was included as a covariate in the model to account for overall differences in response time between subjects, which were not an effect of interest. Rather, we were interested in relative effects within subjects due to the parameters of the cue and S2. Only participants that completed both experiments were included in the ANOVA (N = 9). All correct responses (‘hits’) were included in the analysis for a total of 7395 observations (4308 hits in the WEAK condition and 3087 hits in the STRONG condition) from a set of 18122 total trials (10613 in the WEAK condition and 7509 in the STRONG condition). An alpha criterion of $p \leq 0.05$ was used to define significance.

RESULTS

Group-level response time data are illustrated in **Figure 7**. In the following section we report estimated marginal means and standard errors for the general linear model of the ANOVA. These means reflect the data after controlling for interindividual variability in overall response time, which aids in the interpretation of the interactions of interest.

A main effect was found only for the factor of ‘subject’ ($F_{(1,7382)}=9.425$, $p<0.01$), which was a nuisance variable in this analysis. There was a trend toward significance on the factor of ‘validity’ ($F_{(2,7382)}=5.072$, $p=0.089$). Overall, RTs following valid cues (estimated marginal mean ± SEM: 629.2±3.1ms) were faster than those following neutral cues (653.9±5.5ms), whereas RTs following invalid cues (694.4±6.2ms) were slower than those following neutral cues. RTs did not differ between STRONG and WEAK stimuli (674.1±4.5ms and 644.2±3.8ms, respectively; $F_{(1,7382)}=1.339$, $p=0.411$). RTs did not differ between color and motion targets (657.8±4.2ms and 660±4.2ms, respectively; $F_{(1,7382)}=0.010$, $p=0.932$), consistent with our aim to match task difficulty across the two features.

The ‘validity’ X ‘object’ interaction was significant ($F_{(2,7382)}=18.945$, $p=0.050$). RTs were faster following valid cues than neutral cues, but only when S2’s were WEAK (605.2±4.0ms and 641.6±7.2ms for valid and neutral cues, respectively). When S2’s were STRONG, RTs following valid cues (653.2±4.7ms) did not differ from those following neutral cues (666.1±8.4ms). RTs to validly-cued STRONG S2’s were significantly slower than those to validly-cued WEAK S2’s. RTs following invalid cues were slower than those following neutral cues for both STRONG (invalid: 703.1±9.4ms) and WEAK (invalid: 685.7±8.0ms) S2’s.

The ‘validity’ X ‘feature’ interaction was significant ($F_{(2,7382)}$=31.893, p=0.031). RTs to validly-cued color targets (620.1±4.4ms) were faster than those to validly-cued motion targets (638.3±4.2ms).

The ‘feature’ X ‘object’ interaction was significant ($F_{(1,7382)}$=59.827, p<0.01). RTs to WEAK motion targets (633.6±5.2ms) were faster than those to WEAK color targets (654.8±5.6ms), whereas response times to STRONG motion targets (687.4±6.5ms) were slower than those to STRONG color targets (660.9±6.2ms).

The three-way interaction (‘validity’ X ‘feature’ X ‘object’) was not significant ($F_{(2,7382)}$=0.184, p=0.832).

DISCUSSION

We found that attention could be successfully biased between the visual features of color and motion as indexed by response time effects. That is, the classic cost-benefit effects of invalidly versus validly cuing an upcoming to-be-attended feature dimension were observed here, with participants showing slowed reaction times to invalidly cued targets and speeded reaction times when the cue accurately predicted the feature dimension of the impending judgment. However, this ability to selectively attend was modulated by the motion type of the S2, which was specifically manipulated to differentially engage binding processes. The key finding here is that the deployment of anticipatory feature-based attention was considerably more effective when the invocation of automatic binding processes was weak.

To recall, we made three specific predictions, each of which was confirmed. Firstly, we predicted that RTs for validly-cued WEAK stimuli would be faster than RTs for neutrally-cued WEAK stimuli (i.e., a response time benefit for valid cues), which was indeed the case. Secondly, we predicted that STRONG stimuli would show a smaller response time benefit. In fact, STRONG stimuli did not show a benefit for valid cues at all. Finally, we predicted that response time *costs* for invalid cues would not differ between WEAK and STRONG conditions, because attention was equivalently biased against the target feature in each case. This prediction was also confirmed.

Physiological implications

The interference effects illustrated by this study suggest that feature-based selective attention and binding for object perception are processes that share at least one common anatomical substrate or physiological mechanism. One candidate mechanism that could be common to both of these processes is alpha-band oscillations. Changes in alpha band power have been observed for manipulations of both feature-based attention and object perception.

We have previously reported a role for alpha band oscillations in feature-based selective attention (**Chapter 1**; Snyder and Foxe, 2010). Specifically, alpha-band power increased during preparatory intervals in brain regions linked to the processing of the potentially distracting feature, a pattern of results consistent with the suppressive role of alpha-band power increases reported in spatial selective attention studies (e.g., Kelly et al., 2006; Thut et al., 2006; Worden et al., 2000).

In an MEG study of object processing, Vanni and colleagues (1997) reported greater *decreases* in alpha band power following the presentation of intact line drawings compared to their scrambled and incoherent counterparts. The researchers linked this alpha band effect to object perceptual processes (as opposed to low-level stimulus effects) by asking subjects to report after each trial whether they perceived a coherent object or not. Short (30-106ms) durations were used such that subjects typically only correctly reported perceiving coherent objects 50% of the time. Alpha band responses to misidentified objects were intermediate to those of correctly identified objects and correctly identified scrambled images, indicating that alpha-band decreases relate to the perception of the stimulus as an object, rather than its stimulus attributes.

Alpha band oscillations have therefore been implicated in both feature-based selective attention and coherent object perception, making them a promising candidate for further study of the physiological substrates underlying the interplay between attention and object perception processes. It is noteworthy that alpha-band activity emanates predominantly from the right parietal lobe; precisely the regions implicated in both attentional deployment (Corbetta, Kincade, Ollinger, McAvoy and Shulman, 2000; Foxe, McCourt and Javitt, 2003; Heilman and Van Den Abell, 1980; Vallar and Perani, 1986) and in the attentional coordination of binding (Eglin et al., 1989, 1991, 1994; Esterman et al., 2000; Laeng et al., 2002; Pavlovskaya et al., 2002). Other potential avenues of investigation include other oscillatory frequency bands. In particular, object perceptual processes have been linked to the beta (Sehatpour et al., 2008) and gamma (Eckhorn et al., 1988; Gray, König, Engel and Singer, 1989; Tallon-Baudry, Bertrand, Delpuech and Pernier, 1996) bands. Gamma band activity has also been linked to feature-based selective attention (Keil and Müller, 2010).

Alternative interpretations

Because we administered STRONG and WEAK stimuli in separate sessions, it is possible that participants' knowledge of the nature of the stimuli influenced their preparatory approach. If participants did not bias attention to the same extent in anticipation of a STRONG stimulus as they did in anticipating a WEAK one, this could result in decreased RT benefits for valid cues. However, this would also be expected to result in decreased RT *costs* for *invalid* cues, which was clearly not what was found. This suggests that the preparatory approach was the same for both classes of stimuli, and that decreased biasing due to binding demands manifests during stimulus-processing stages.

From a theoretical standpoint, it is unlikely that the quality of processing of the relevant feature is reduced due to increased binding demands. Indeed, this would be contrary to the literature regarding the spread of attention within an object, which has hitherto reported only enhancement effects (e.g., O'Craven et al., 1999; Schoenfeld et al., 2003; Wylie et al., 2004; Molholm et al., 2007; Katzner et al., 2009; Fiebelkorn et al., 2010a, 2010b). Furthermore, if the quality of processing of the relevant feature was reduced, RTs following neutral cues would be predicted to be slower for STRONG stimuli compared to WEAK stimuli, and this was not observed here. Thus we conclude that the quality of processing of the irrelevant stimulus feature was likely enhanced during STRONG object perception, resulting in a decreased biasing between the two competing features. The key implication of this is that a certain "quality threshold" of processing must be met for features to be bound together, and that this requirement dominates prior attentional deployments during object perception.

Methodological considerations

A critical feature of the current study is that participants' performance levels were held below ceiling, indicating that processing capacity was exhausted. Tasks performed at ceiling levels may leave processing resources to spare such that the functions under investigation do not have to compete, obviating the need for attentional biasing. That is, it may be the case that enhanced processing of irrelevant features in a selected object is a convenient thing to do with excess processing resources, but is not strictly a fundamental part of the object-perception process.

For example, O'Craven and colleagues (1999), in a neuroimaging study, showed participants overlaid images of faces and houses, one of which moved on each trial. In the critical condition, the participants attended to the direction of motion of the moving image, but the content of that image was irrelevant to the task. The key result was that metabolic signatures increased in face-responsive brain areas when the face was moving and increased in house-responsive areas when the house was moving, notwithstanding the task-irrelevance of these object categories. This result was taken to indicate the spread of attention to the task-irrelevant category of the attended image. However, in this case the task was likely quite easy (behavioral results were not reported, but participants merely had to discriminate between 'up', 'down', 'left', and 'right'), and there would be no disadvantage conferred by allocating excess attentional resources to processing the object category. In fact, it may be the case that parsing the object category assisted the participants in segregating the moving image from the spatially colocalized static image. Thus, one cannot conclude from these results that the spread of attention to irrelevant features is a fundamental process rather than a merely convenient one.

To test the robustness of the cross-feature attentional-spreading phenomenon, it is necessary to create a context for which it would be advantageous for attention *not* to spread between the features of an object (i.e., the two features are competitors for the same processing resources) as in the current study. In such a context, evidence of attentional spread (such as decreased biasing of attention between the shared features of an object) would indicate the automaticity of the process. To our knowledge, the current study is the first to demonstrate that cross-feature spread of attention occurs even when explicitly detrimental to task performance.

Other findings

In addition to the predicted interaction between cue validity and object strength, we found two-way interactions between which feature was cued and the effects of cue validity and object strength. These interactions could be related to the temporally integrative nature of motion judgments. For example, we found that response times to validly cued motion targets were delayed 18ms compared to validly cued color targets. This delay might reflect the fact that motion direction required at least two video frames to determine, whereas color information was available in a single frame. Likewise, STRONG motion targets, which had curved motion, required at least three frames to resolve (i.e, two video frames to establish direction, and then an additional frame for the change in direction). It would not be surprising if the brain required more than this minimum number of frames to assess the motion curvature in this demanding task. This may account for why this condition had the greatest delay in response time compared to other target types.

Conclusion

We found that response time benefits for valid feature-based attentional cues that were observed when imperative stimuli only weakly engaged binding processes were reduced when imperative stimuli strongly engaged binding. This result suggests that feature-based selective attention processes and binding processes for coherent object perception rely at least partially on a common physiological substrate. One potential candidate substrate is alpha band oscillations, which have been separately implicated in both processes.

Figure 6 (caption): Schematized illustration of target examples. Arrows represent motion and were not actually present in the stimulus. The length of each arrow represents the speed of the adjacent dot. For both WEAK and STRONG stimuli, color targets were defined by the presence of two colors of dots while the dots continued to move on a common linear path. For WEAK stimuli, motion targets were defined by the simultaneous presence of two different motion directions with uniformly colored dots. For STRONG stimuli, motion targets were defined by the presence of curved motion (i.e., sequential presentation of different motion directions). Note that no one particular color or motion direction was indicative of target presence. Stimuli are shown here on a white background for illustration; stimuli were on a black background in the experiment. Color and motion differences are enhanced for clarity. Dots in the actual stimuli were smaller and more numerous.

Figure 6:

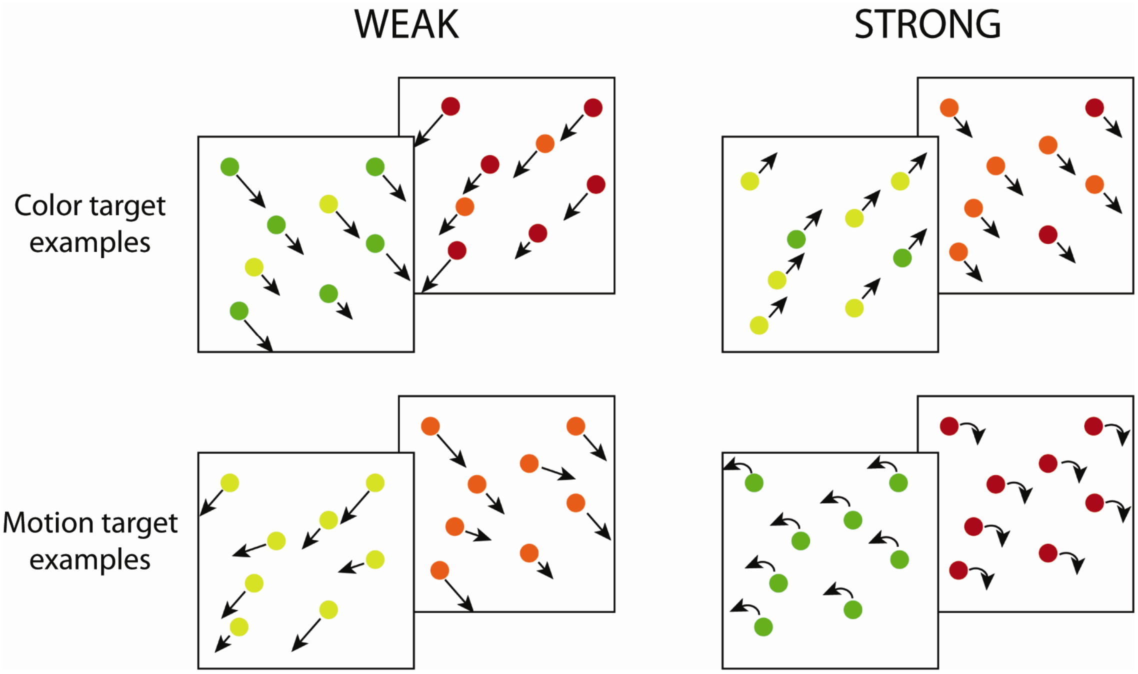

Figure 7 (caption): Summary of behavioral results. For each target feature, color or motion, response times (RTs) are summarized for each type of cue (valid, neutral, or invalid). For WEAK objects, RTs were faster for validly-cued stimuli compared to those for neutrally-cued stimuli. For STRONG objects, however, the RTs for validly-cued stimuli did not differ from those for neutrally-cued stimuli.

Figure 7:

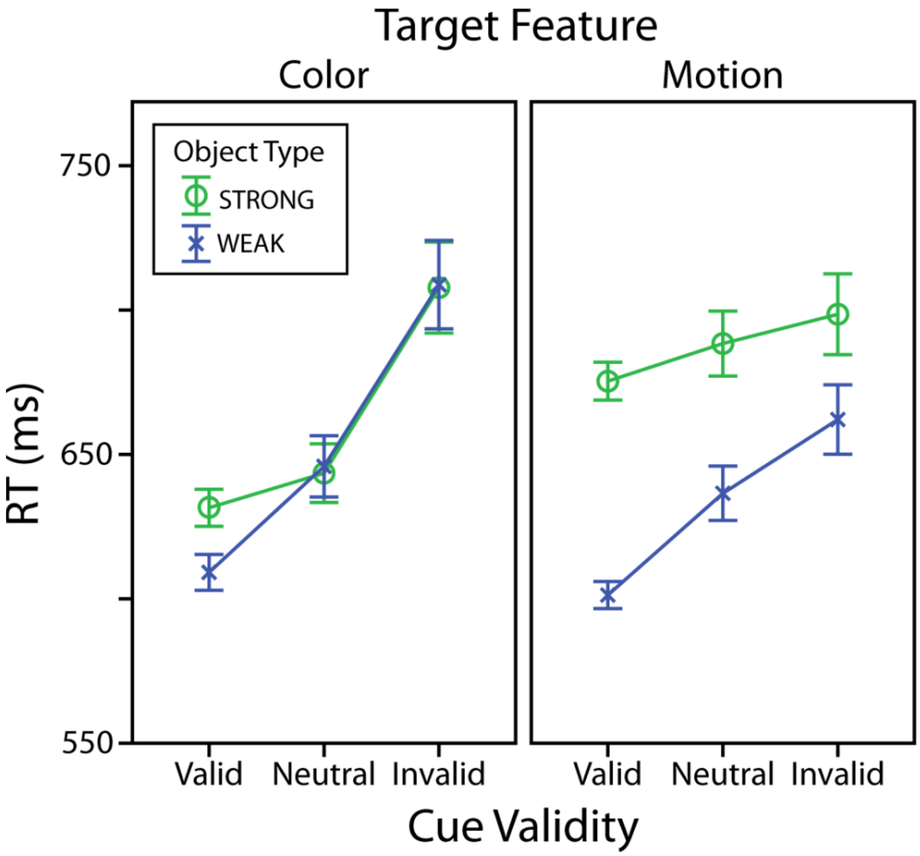
Target Feature
Color
Motion
Object Type
STRONG
WEAK
RT (ms)
750
650
550
Valid
Neutral
Invalid
Valid
Neutral
Invalid
Cue Validity

CHAPTER 3

Pitting Binding Against Selection: Electrophysiological Measures of Feature-Based Attention are Attenuated and Delayed by Gestalt Object Grouping.

It has been estimated that approximately 10^7 bits of data travel down the optic nerve each waking second (Koch et al., 2006), and while our perception is impressive and creates an illusion of continuity, we consciously register but a fraction of that information. Each day we must extract useful signals from the confusion of information in our environment if we are to perceive, learn, remember, and plan actions useful to our survival. This filtering of information from the environment is achieved through a combination of endogenous (“top-down”) and exogenous (“bottom-up”) biasing that resolves the competition for limited processing resources. The top-down form of this biasing, known as “selective attention”, involves allocating processing resources such that sources or classes of information expected to be relevant in a given context are given priority at the expense of irrelevant or distracting information (e.g. Broadbent, 1958, Kahneman, 1973; Neely, 1977; Schneider and Fisk, 1982; Schneider and Shiffrin, 1977; Shiffrin and Schneider, 1977; Foxe et al., 2005a). In this chapter we focus on selective attention to basic visual features, such as color and motion. Feature-based selective attention is an efficient way to guide visual search, as when meeting a friend at a crowded train station who has mentioned in advance, “I’ll be wearing red.” Experimentally, people informed in this way about the likely color of an upcoming task-relevant stimulus are faster to respond if the stimulus does in fact appear with the expected color than if it appears with another color (e.g., Egner et al., 2008; Most and Astur, 2007. See also Hopf, Boelmans, Schoenfeld, Luck and Heinze, 2004). Moreover, neurophysiological indices of feature-processing are also enhanced for stimuli with an attended feature, regardless of spatial location (e.g., Corbetta et al., 1991; Martinez-Trujillo and Treue, 2004).

Seemingly in opposition to this intentional boost in processing of separable features is an inherent bias to integrate features that co-occur in time and space. Previous literature has demonstrated that enhanced processing spreads from an object's task-relevant features to its irrelevant features (Albrect et al., 2008; Egly et al., 1994; Fiebelkorn et al., 2010a, 2010b; Katzner et al., 2009; Molholm et al., 2007; O’Craven et al., 1999; Schoenfeld et al., 2003; Wylie et al., 2004). What's more, this spreading of attention across the features of an object seems to be highly automatic (i.e., "bottom-up"). In a prior experiment by our group (Fiebelkorn et al., 2010a), which investigated the spread of attention from visual to auditory features of well-known multisensory objects (dogs, cars and guitars), we found a spread of attention to task-irrelevant auditory stimuli that were semantically incongruent with a co-occurring visual stimulus (e.g., the sound of a car paired with a picture of a guitar when dogs were attended). In other words, auditory stimuli received a processing enhancement merely for co-occurring with a visual stimulus during a visual task, even though attention to the auditory stimuli was potentially detrimental to task performance (due to incongruent pairings).

In the present experiment we placed feature-based selective attention and bottom-up binding processes in direct conflict to test whether top-down biasing can supersede the object-based spreading of attention under conditions where grouping of basic features is detrimental to task performance. To this end, we used a cued attention task to direct attention to one of two basic visual features (color or motion). Participants then made a difficult discrimination with respect to the cued feature (**Figure 8**). At the same time, we used well-established Gestalt grouping principles to manipulate the degree to which the imperative stimulus was perceived as a single coherent object (i.e., we differentially engaged binding processes; see **Chapter 2**, also Snyder and Foxe, 2011). This manipulation of Gestalt grouping was irrelevant to the behavioral task, and many participants reported an unawareness of the grouping manipulation upon debriefing. The optimal approach for our participants would have been to suppress the distracting (i.e., task-irrelevant) feature, which would require that feature-based selective attention disrupt or diminish bottom-up binding processes that typically occur in response to a coherent object. To test whether feature-based selective attention effects indeed persist despite environmental cues that would typically lead to an object-based spreading of attention, we examined event-related potentials (ERPs) recorded in response to dot arrays, where either a color or a motion task was cued. These ERPs were compared across conditions for which we manipulated the degree to which the dot arrays would be perceived as a single object or an array of relatively independent elements, which we termed the STRONG and WEAK object conditions, respectively. In the STRONG condition, all dots moved at the same speed and in the same direction, which leads to the percept that the dots are ‘painted’ on a single transparent surface that is itself sliding past an aperture. In the WEAK condition, dots moved in the same direction, but each at a unique rate, which does not lead to the percept of a single object.

METHODS

Participants

Sixteen adults (11 male, 3 left-handed) aged 20 to 47 years (mean±s.d.: 27.19 ± 6.75 y) participated in the experiment. Three participants were excluded due to data quality concerns. The resultant sample contained 10 males and 3 females aged 20 to 47 years (27.92 ± 7.32y), three of whom were left-handed. Participants were sourced from the undergraduate and graduate student populations of The City College of New York, and from the local community. None of the participants had any history of brain injury or disease, per self-report. Participants had normal or corrected-to-normal visual acuity and normal color vision per self-report. All participants provided informed consent prior to the experiment. All materials and procedures were approved by the institutional review board of The City College of New York in accordance with the United States Public Health Service Act (US 45 CFR 46) and the Declaration of Helsinki.

Stimuli

The experiment was administered in a light- and sound-attenuated chamber using Presentation software version 14.4 (Neurobehavioral Systems, Albany, California). All stimuli were presented on a 34.5 x 55.0 cm LCD monitor with a 60 Hz refresh rate (ViewSonic, model VP2655wb). Trials began with a warning cue consisting of a white fixation dot on a black background for 1 s, followed by a cue word (‘COLOR’, ‘HUE’, ‘MOTION’, or ‘DIRECTION’) in white block capitals for 1 s. The words ‘color’ and ‘hue’ both directed attention to the color of the upcoming imperative stimulus (the second stimulus, or “S2”). Likewise, ‘motion’ and ‘direction’ both directed attention to the motion of the stimulus. The use of multiple cue words for each feature was to reduce the automatization of the task through implicit learning, with an aim to maintain the engagement of endogenous orienting mechanisms throughout the session. After an interval of 1.7 to 2.3 s (random, and evenly distributed) during which only a black screen was displayed, the S2 was presented for 0.2 s. Each S2 was followed by a 1 s response interval. Each subsequent trial began immediately following the response interval (**Figure 1**).

The S2 consisted of an array of one thousand dots, each subtending 0.05 degrees of visual angle, constrained to a square aperture subtending 5 degrees of visual angle. In one condition, designed to engage Gestalt grouping processes, all dots moved in the same direction at a speed of 21 degrees of visual angle per second. We refer to this condition as the STRONG condition, indicating that Gestalt processes are strongly engaged. A common percept of this type of stimulus is that the dots are 'painted' on a single transparent surface that is itself sliding past the aperture. In another condition, designed to minimize the engagement of Gestalt grouping processes, each dot moved at a unique speed between 14 and 28 degrees of visual angle per second, and all dots moved in the same direction. For this type of stimulus the dots are not perceived as parts of a single object as is the case when the dots all move at the same speed. We refer to this condition as the WEAK condition, indicating that Gestalt processes are weakly engaged (relative to the STRONG condition). Dots 'wrapped around' the edges of the square aperture, so that the total amount of illumination was held constant. STRONG and WEAK S2's each occurred on 50% of trials in an unpredictable order. It should be emphasized that the STRONG/WEAK manipulation was irrelevant to the task. Participants were not informed by the experimenter of this manipulation until debriefing, at which point several participants expressed that they were unaware of the STRONG/WEAK distinction during the experiment.

Dots were colored with hues from an isoluminant plane of DKL color-space (Derrington et al., 1984). This color-space uses the response properties of neurons in macaque lateral geniculate nucleus to create a subjective luminance axis, planes orthogonal to which are approximately isoluminant. The use of this color-space enables the continuous variation of hue needed to derive hue discrimination thresholds while controlling for subjective luminance.

Task

Two types of coloration patterns and motion types were used for the S2's. Half of the trials consisted of uniformly-colored S2's (color standard), whereas for the other half of trials 20% of dots had a slightly different color from the majority (color target). Orthogonally to this color manipulation, half of trials contained S2's for which all dots moved on a common linear trajectory (motion standard) and the other half contained S2's for which the dots moved on a curved trajectory (motion target). Schematized examples of targets are illustrated in **Figure 8**. The degree of difference for each of the relevant targets was titrated on a per-subject basis to 80% detection rate prior to beginning each experiment using an up-down transformed response (UDTR) modified staircase procedure (Wetherill and Levitt, 1965). No particular value of any feature indicated a target: subjects had to detect a particular feature *variation* in the stimulus. This strategy was used to reduce competition within a feature processing area (if subjects were attending to red and suppressing green, for example). The goal, rather, was to have subjects attend to *color* and suppress *motion*, or vice versa as the cue indicated.

Participants were instructed to attend to the cued feature only, ignoring the uncued feature entirely, and to respond with a button press upon detection of a target in the cued feature. The requirement to ignore the uncued feature was intended to make any target properties appearing in the uncued feature distracting, so that object-related spreading of attention could provide no task-related benefits. Participants completed 10 ten-minute blocks, which included breaks every 12 trials to reduce fatigue and maintain a high level of alertness throughout the session.

EEG Recording

Continuous EEG was acquired through the ActiveTwo BioSemi (Amsterdam) electrode system from 168 scalp electrodes, digitized at 512Hz. With the BioSemi system, every electrode or combination of electrodes can be assigned as the reference, which is done purely in software after acquisition. BioSemi replaces the ground electrodes that are used in conventional systems with two separate electrodes: Common Mode Sense (CMS) active electrode and Driven Right Leg (DRL) passive electrode. These two electrodes form a feedback loop, which drives the average potential of the subject as close as possible to the reference voltage of the analog-to-digital converter, thus rendering them references. One electrode placed 1 cm posterior to each orbital canthus and one electrode placed on the nasion were used to monitor eye movements.

Data processing

EEG data were processed using the FieldTrip toolbox (Donders Institute for Brain, Cognition and Behaviour, Radboud University Nijmegen, the Netherlands) for MATLAB (The MathWorks Inc., Natick, Massachusetts). Only trials with correct behavior (responses to targets and withheld responses to non-targets) were included in the analysis. Data were epoched from 100ms pre-S2 onset to 500ms post-S2 onset. Each trial was first visually inspected, and channels showing transient, large-amplitude electronic artifacts were linearly interpolated on a per-trial basis. Trials with more than 5 interpolated channels were excluded from further analysis. Then, the data were decomposed into independent components that were used for correction of muscular and ocular artifacts.

To perform the independent components analysis (ICA), we first concatenated all the trials for each individual participant. This resulted in very large datasets, and so we next downsampled the data by a factor of 4 for efficiency of the ICA algorithm. The concatenated, downsampled data were then used to derive unmixing matrices using the FastICA algorithm with default settings (Hyvärinen, 1999). The downsampled data were then discarded, and the unmixing matrices were used to unmix the original, undownsampled data for each participant into independent components. Components reflecting artifacts were identified using an automated approach described by Whittingstall, Bartels, Singh, Kwon and Logothetis (2010). Artifacts were identified based on temporal and spatial characteristics. For the temporal criterion, component time courses were rectified and then converted to z-scores, and components containing z-scores exceeding a value of 7 were discarded. This criterion identifies infrequently-occurring, large amplitude, transient artifacts. For the spatial criterion, each component topography (i.e., each column of the unmixing matrix) was converted to z-scores, and components for which the largest z-score exceeded a value of 9, or for which the sum of the two largest z-scores exceeded 15, were discarded. This criterion identifies artifacts that are limited to one or two electrodes. The choice of z-score threshold values was adopted from the report by Whittingstall and others (2010). In a preliminary stage, we tested this method with several z-score thresholds and found that the particular set of components rejected was consistent across a moderate range of threshold values. After rejecting components reflecting artifacts, the remaining components were remixed into sensor-space and re-divided into individual epochs for further processing and analysis.

Artifact-corrected data were re-referenced to a frontopolar electrode equivalent to Fpz in the 10-20 system. This reference was chosen because it was a clean recording site in all participants far from occipital cortices, which were hypothesized to be the regions likely to show maximal effects in this visual paradigm. Data were then low-pass filtered below 30 Hz and then high-pass filtered above 0.3 Hz. In each case a 6th-order zero-phase digital Butterworth filter was used. Expanded epochs (2s pre-S2 to 2s post-S2) were used for filtering to avoid contamination by artifacts introduced by discontinuities at the edge of the epochs. Subsequent to filtering, the epochs were trimmed back to 100ms pre-S2 onset to 500ms post-S2 onset and averaged across trials. A mean of 161 trials (± 25.4 s.d.) were included in the ERP for each condition for each participant.

Analysis

Behavioral Analysis: We analyzed both performance accuracy and response times. In each case we used a repeated-measures ANOVA with factors of ‘feature’ (two levels: color and motion) and ‘S2 type’ (two levels: WEAK and STRONG). Our dependent measure for performance accuracy was the discriminability index (d’), a measure of performance which takes into account response bias. For response time (RT), we tested the median RT’s across trials for each participant for each cell of the ANOVA design matrix. We chose the median as the measure of central tendency because it is robust to outliers, which were likely in this case because participants were not instructed to respond rapidly. Behavioral statistics were analyzed using PASW version 18.0.0 (SPSS Inc., Somers, New York).

Statistical Hypothesis Testing of ERP Differences: We had a directed hypothesis that the ERP indices of attentional biasing would differ between STRONG and WEAK S2s. To identify the ERP indices of our feature-based attention manipulation, we performed a running dependent samples *t*-test between conditions of attention-to-motion and attention-to-color for all channels and time points. We used a clustering approach to control for inflation of type I error due to multiple comparisons (Guthrie and Buchwald, 1991). The rationale for this method is that type I errors are unlikely to occur simultaneously at adjacent electrodes and unlikely to endure for several consecutive time points. Since the EEG signal does not change arbitrarily fast, however, there is some dependence between consecutive time points. We examined the autocorrelation of the noise in the baseline interval of the ERP to determine the minimum time delay at which time points in the EEG are not significantly correlated, which we found to be 15 samples (29ms). We thus required two-tailed *p* values below 0.05 to be simultaneously present at two or more adjacent electrodes and to persist for at least 15 samples to consider the effects significant. This statistical analysis was applied separately to the STRONG and

WEAK conditions, and then the STRONG and WEAK attention effects were compared. We also performed a running dependent samples *t*-test comparing all STRONG to WEAK responses, regardless of which feature was attended, to characterize main effects of the Gestalt grouping manipulation.

Source Estimation: To visualize the cortical regions most likely to be contributing to effects observed at the scalp, we employed a minimum norm current source estimation (MNE) method using Brain Electrical Source Analysis software (BESA GmbH, ver. 5.1.8), using the following settings: depth weighting and spatiotemporal weighting using subspace correlation (dimension = 7); and noise estimation using the 15% lowest values, weighted by averaging over channels. Source localization was applied to subtraction waveforms representing effects of interest. Subtractions were made as follows: 1) attend-color WEAK minus attend-motion WEAK; 2) attend-color STRONG minus attend-motion STRONG; and 3) STRONG minus WEAK.

RESULTS

Behavior

Performance Accuracy: Performance results are summarized in **Figure 9a**. Discriminability scores ranged from d'=0.88 to d'=4.22 (mean±s.d.: 2.44 ± 0.76), indicating that all participants performed above chance (i.e., d'>0) and below ceiling (i.e., d'<4.65, which reflects a 99% hit rate and 1% false alarm rate). Results of the ANOVA revealed no significant main effects or interactions for the factors of 'feature' and 'S2 type' on *d'* scores (all *p*'s >0.12;). The lack of significant differences was consistent with our aim to equate task difficulty across conditions for each subject.

Response Times (RT's): Response time results are summarized in **Figure 9b**. Median RT's ranged from 511.7 ms to 957.0 ms (mean±s.d.: 698.8 ± 113.9 ms). Results of the ANOVA revealed a main effect of 'feature' ($F_{1,12} = 6.593$, $p = 0.025$). A post-hoc paired-samples *t*-test indicated that responses to motion targets were significantly slower than responses to color targets ($p = 0.025$). On average, participants responded 42.2 ms (s.d.: 59.3 ms) slower to motion targets than to color targets, and 10 out of 13 participants showed slower RTs for motion than for color. Slower RT's for motion judgments compared to color judgments are not uncommon in feature-based attention studies (e.g., Anllo-Vento and Hillyard, 1996; Schoenfeld et al., 2007), and likely reflect the inherent temporally-integrative nature of motion processing.

No main effect was found for the factor of 'S2 type' ($F_{1,12} = 1.241$, $p = 0.287$). The 'feature' by 'S2 type' interaction was not significant ($F_{1,12} = 0.207$, $p = 0.657$).

Statistical Testing of ERP Differences

For WEAK stimuli, significant differences between attention-to-color and attention-to-motion were found at parieto-occipital electrodes from 180 ms to 221 ms post-stimulus onset (**Figures 10, 11a**). The ERP was significantly more negative for attention-to-color than for attention-to-motion at these times and locations. In contrast, no attention effects for STRONG stimuli were found (**Fig. 10**).

A comparison between the ERP's for STRONG and WEAK stimuli revealed a robust effect at parieto-occipital electrodes from 273 ms to 404 ms (**Fig. 12a,b**). The ERP was more positive for STRONG stimuli than for WEAK stimuli at these times and locations.

For both WEAK and STRONG stimuli, no effects were observed at any time point for electrodes located over the eye muscles, suggesting that eye movements were not a confounding factor.

Source Estimation

The MNE analysis revealed that the early (180-221 ms) attention effect for WEAK stimuli was attributable to current sources in occipital pole regions and ventral occipito-temporal cortex bilaterally (**Figure 11b**). The source estimation for the general Gestalt binding effect (i.e., STRONG vs. WEAK) implicated right lateral occipital cortex and left dorsolateral prefrontal cortex (**Figure 12c**).

DISCUSSION

Here we tested whether the electrophysiological effects of feature-based selective attention persist despite environmental cues (i.e., common fate) that typically lead to object-based binding (Albrecht et al., 2008; Busse, Roberts, Crist, Weissman and Woldorff, 2010; Egly et al., 1994; Fiebelkorn et al., 2010a; O'Craven et al., 1999; Schoenfeld et al., 2003). In other words, we examined whether the well-established ability to bias processing toward a specific feature (color or motion) would be counteracted by the well-established bias to form coherent objects. To test the automaticity of the spreading of attention, which is thought to reflect feature binding, we designed a task where object-based integration would be detrimental to performance.

On each trial, participants' were cued to detect either color or motion targets, while overall performance was held at approximately 80 percent. Pinning performance at less than 100 percent accuracy served to saturate attentional resources, such that any boost in the processing of the uncued (or irrelevant) feature would necessarily detract from processing of the cued (or relevant) feature. If Gestalt principles in the STRONG condition led to an object-based spreading of attention from the cued feature to the uncued feature, we would predict that the two features would be processed with greater parity. Evidence of decreased feature-based selectivity in the presence of environmental cues for object-based binding would suggest that the inherent bias to process objects as wholes supersedes task-driven attempts to bias processing toward a particular feature in order to optimize performance.

Our electrophysiological measure of feature-selective processing was the difference between ERPs when color was attended compared to when motion was attended. For the WEAK condition, robust evidence of biased processing for the cued feature was evident at approximately 180 ms; whereas, for the STRONG condition, evidence of biased processing was not observed. We therefore conclude that feature-based selective attention is substantially attenuated and delayed by object-based binding processes. These data emphasize the strength of the tendency to integrate features that share a common fate, even when doing so will interfere with the task at hand.

The significant differences that we observed between attention to color and attention to motion for the WEAK condition are similar in terms of latency and topography to previously observed feature-specific attention effects (Anllo-Vento and Hillyard, 1996; Hopf et al., 2004; Schoenfeld et al., 2007). Source localizations implicated ventral occipito-temporal regions bilaterally, with additional contributions from occipital pole regions. Because color processing has been previously localized to the ventral occipito-temporal regions (Corbetta et al., 1991; Vaina, 1994), we interpret this pattern of results as enhanced processing in color-related cortical areas when color is attended compared to when motion is attended.

Attention effects are typically measured by comparing the ERPs in response to attended and unattended stimuli, which results in a difference wave often referred to as a "selection negativity" (e.g., Anllo-Vento and Hillyard, 1996). In the present experiment, however, we used a somewhat novel comparison to quantify feature-based attention. We compared the response to physically identical stimuli when different constituent features were attended. This approach was necessary to observe not only the effects of feature-based attention, but also the object-based spread of attention, which is hypothesized to occur within objects for which at least one feature is attended. Despite the fact that all of the stimuli included an attended feature (either color or motion), we predicted that ERP differences at the level of the scalp would occur as a result of the spatial differences in activated processing areas (i.e., dorsal and ventral visual regions for motion and color processing, respectively; Ungerleider and Mishkin, 1982). Our dependent measure for biased processing is thus not equivalent to a stereotypical selection negativity, despite a resemblance in timing and topography. Rather than an additive component, our effects can be better conceived as reflecting a change in the *distribution* of processing across all the brain areas involved.

Because the evidence of feature-based selective attention was present for WEAK stimuli at 180 ms, but absent for STRONG stimuli at that time, we can conclude that object-based binding effects within STRONG stimuli must have onset prior to 180 ms. An obvious question is how the timing of our current effects accord with timing of known binding processes. Evidence for the timing of object-based binding comes from several electrophysiological investigations of the canonical “illusory contour” stimulus class. These studies have utilized Kanisza illusory figures (cf. Kanisza, 1976), which are composed of arrays of Pacman shaped inducers oriented with their “mouths” aligned toward a common point, such that the region interior to the inducers is perceived as a geometric shape superimposed on the background, although no such shape actually exists. If the inducers are sufficiently misaligned the illusion is destroyed. By comparing the ERP to the Kanisza figures to their non-illusion-inducing counterparts, an “illusory contour effect” has been identified that onsets around 100 ms and peaks around 170 ms (Herrmann, Mecklinger and Pfeifer, 1999; Fiebelkorn et al., 2010b; Foxe, Murray and Javitt, 2005b; Murray, Foxe, Javitt and Foxe, 2004; Murray, Imber, Javitt and Foxe, 2006; Murray et al., 2002; Shpaner, Murray and Foxe, 2009), the effect is taken to indicate binding of the disparate inducers into a single Gestalt, as in the current experiment.

While our inference that binding effects occurred prior to 180ms was indirect, we directly observed effects of our Gestalt grouping manipulation during later time periods (273 ms to 404 ms). This effect was characterized by greater positivity at posterior electrodes, and implicated a robust current source in right lateral occipital cortex. Lateral occipital cortex has previously been implicated in object-related processing (e.g., Lucan, Foxe, Gomez-Ramirez, Satian and Molholm, 2010; Martinez et al., 2006; Martinez, Ramanathan, Foxe, Javitt and Hillyard, 2007b), particularly with regards to integrating fragmented images (Doniger et al., 2000, 2001; Sehatpour et al., 2006) and illusory contours (Martinez, Teder-Sälejärvi and Hillyard, 2007a; Murray et al., 2002, 2004; Shpaner et al., 2009).

Conclusion

We found that participants were successfully able to bias processing resources between the features of color and motion. This biased attentional state was evident in differences in the ERP's between when motion was attended and when color was attended. These effects were relatively early (180ms) and attributed to feature-specific cortical processing regions. However, these differences were only evident when enviromental cues for object-based binding were weak. When environmental cues for binding were strong, evidence for biased feature-specific processing was not observed. We therefore conclude that bottom-up binding processes supercede top-down feature-based selective attention, even when explicitly detrimental to task performance.

While we used a highly-contrived experimental context to reveal the interaction between feature-based selective attention and binding processes, we believe that our results provide insight into how attention and binding processes are used to guide interaction with the environment in a natural setting. We propose the following model: First, feature-based attention, which is not spatially-specific, highlights objects with the attended feature for subsequent spatial selection. Binding processes then take over, which leads to a spread of attention between the attended object's constituent features, and an object-based attentional selection is formed. The attended object can then be scrutinized with respect to the individual features, but this scrutiny is likely deferred to post-perceptual timeframes.

Figure 8 (caption): Schematized illustration of target examples. Arrows represent motion and were not actually present in the stimulus. The length of each arrow represents the speed of the adjacent dot. Color targets were defined by the presence of two colors of dots. Motion targets were defined by the presence of curved motion (i.e., sequential presentation of different motion directions). Note that no one particular color or motion direction was indicative of target presence, and that the examples given here are not exhaustive. Both color and motion targets could be present in a single stimulus. Stimuli are shown here on a white background for illustration; stimuli were on a black background in the experiment. Color and motion differences are enhanced for clarity. Dots in the actual stimuli were smaller and more numerous.

Figure 8:

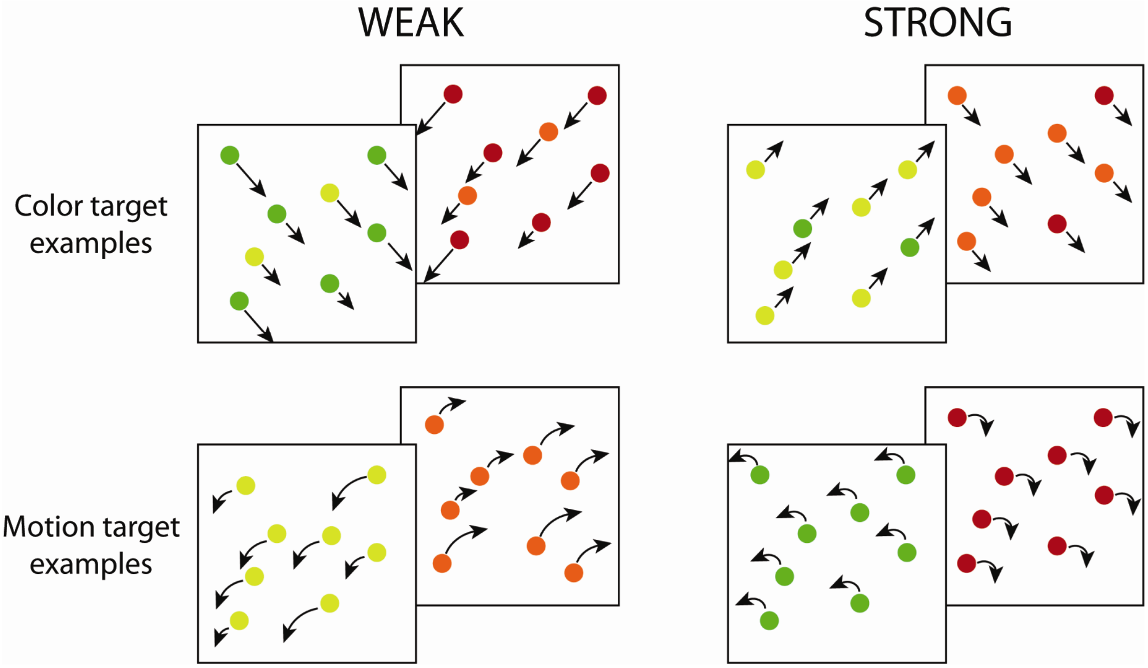

Figure 9 (caption): Summary of behavioral results across the 13 included participants. A) Performance accuracy in d’ scores. B) Response times in milliseconds. ∗: $p < 0.05$ (repeated-measures ANOVA). Error bars indicate 2 standard errors of the mean.

Figure 9:

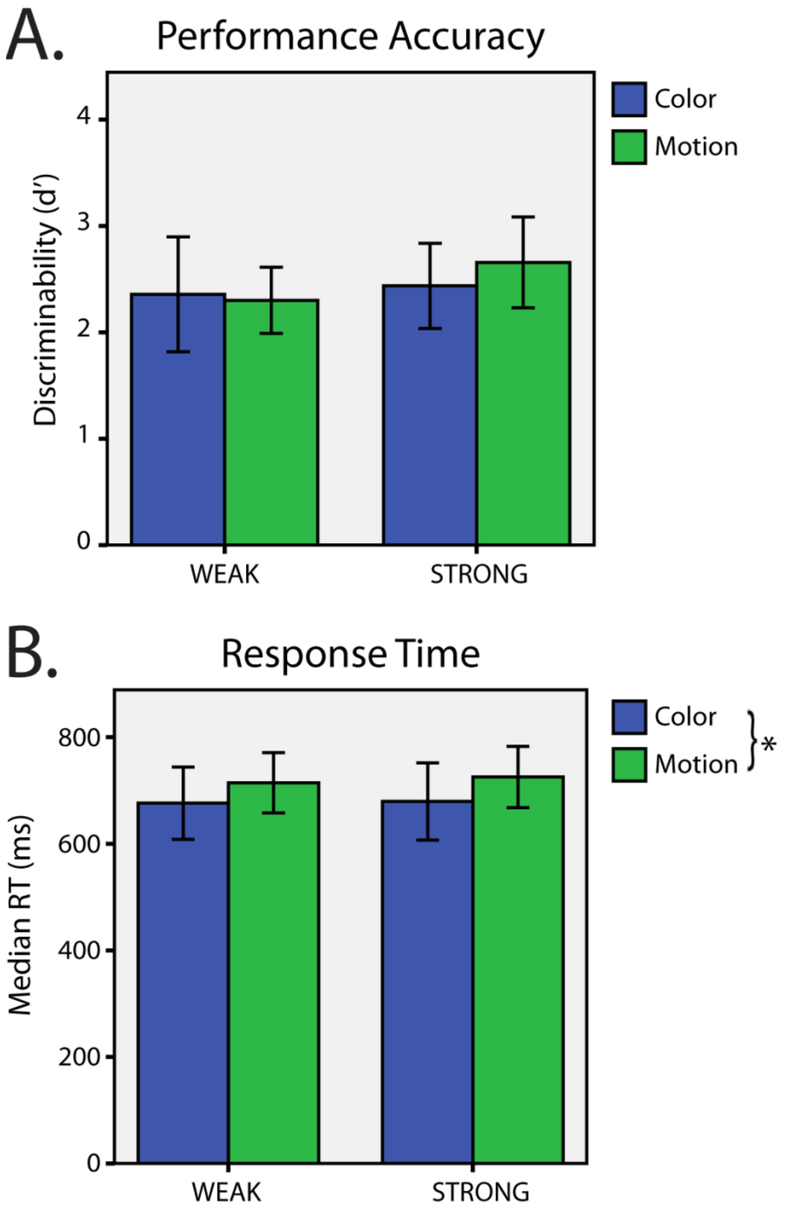
A.
Performance Accuracy
Discriminability (d')
0
1
2
3
4
WEAK
STRONG
Color
Motion
B.
Response Time
Median RT (ms)
0
200
400
600
800
WEAK
STRONG
Color
Motion
*

Figure 10 (caption): Feature-based attention effects for WEAK and STRONG stimuli. ERPs in response to S2's are plotted for a representative electrode (electrode location highlighted in insets) above the p-values for a running t-test of the comparison 'color vs. motion.' The green dashed line represents $p = 0.05$. The light green rectangle highlights the time period of significant difference.

Figure 10:

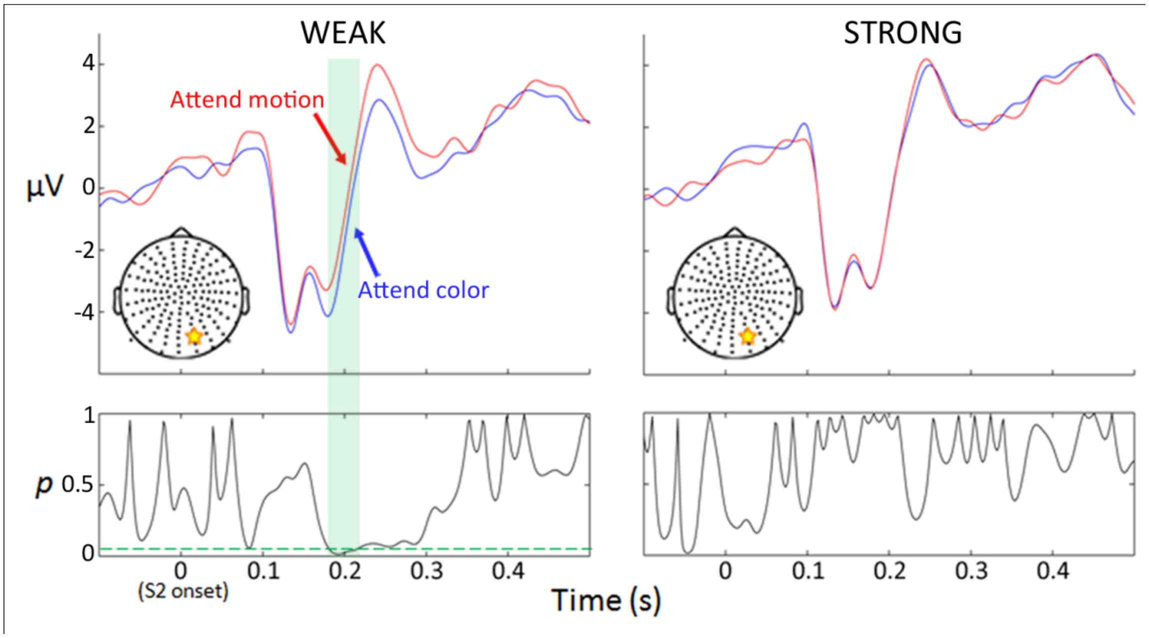
WEAK
STRONG
Attend motion
Attend color
μV
4
2
0
-2
-4
p
1
0.5
0
0
(S2 onset)
0.1
0.2
0.3
0.4
Time (s)

Figure 11 (caption): Maps of feature-based attention effects. A) Scalp topography of attend-color ERP minus attend-motion ERP for WEAK stimuli. Black circles depict significant channels. B) Minimum norm current source estimates for time periods with significant attention effects for WEAK stimuli.

Figure 11:

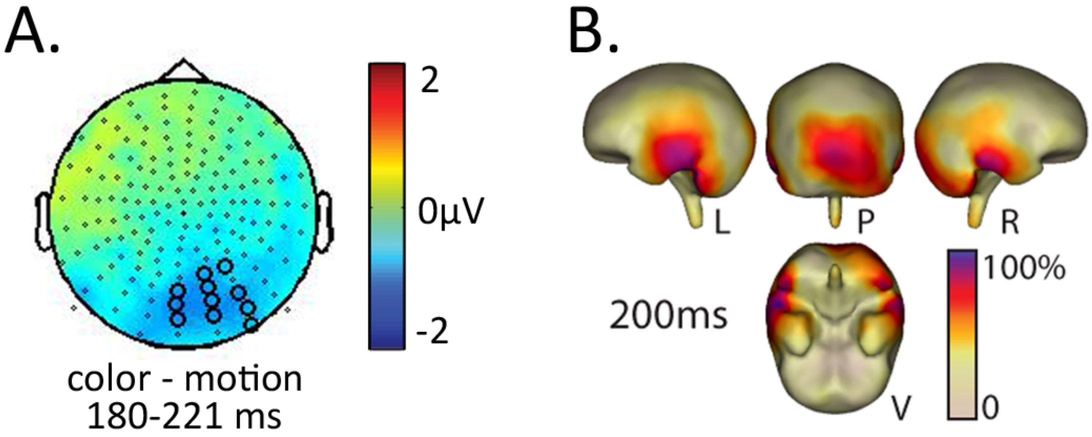

Figure 12 (caption): General Gestalt grouping effect (i.e., STRONG minus WEAK). A) ERP’s for representative electrode (electrode location highlighted in inset). B) Scalp topography. Black circles depict significant electrodes. Green shading depicts the time period of significant differences. C) Minimum norm current source estimate for time period with significant effect.

Figure 12:

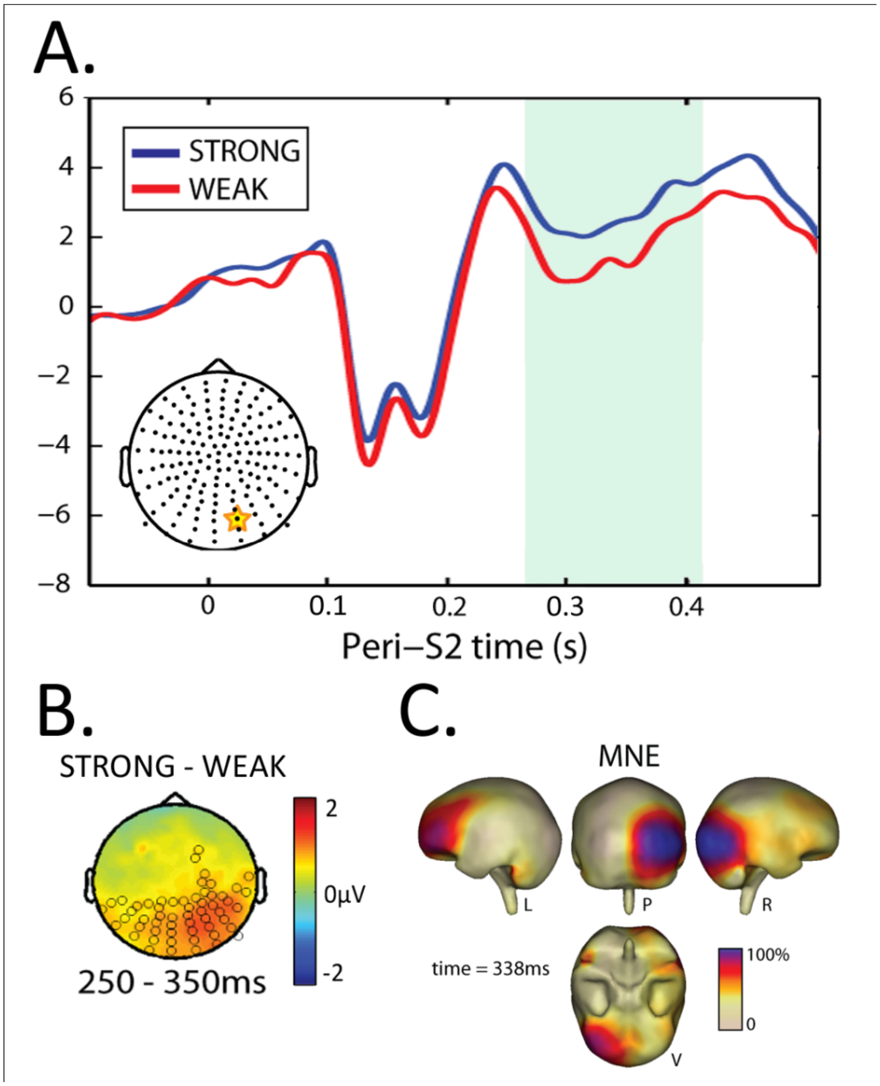
A.
6
4
2
0
−2
−4
−6
−8
STRONG
WEAK
0
0.1
0.2
0.3
0.4
Peri–S2 time (s)
B.
STRONG - WEAK
2
0μV
-2
250 - 350ms
C.
MNE
L
P
R
time = 338ms
100%
0
V

GENERAL DISCUSSION

We have described three sets of experiments that investigated the preparatory mechanisms of visual feature-based selective attention, and the interaction between such top-down attention mechanisms and object-related binding processes. Our results lead us to three key conclusions: (1) anticipatory attentional suppression of visual features invokes alpha-band mediated gating mechanisms, similarly to spatial attention (**Chapter 1**), (2) binding of features for coherent object perception dominates selective attention, even when detrimental to task performance (**Chapter 2**), and (3) such automatic binding processes have already counteracted feature-based attention by 180 ms (**Chapter 3**), although some biased processing is again evident in post-perceptual time frames (i.e., > 500 ms). These findings suggest that feature-based attention is a useful mechanism for guiding visual search, but that once an object is selected by attention, biased processing between features is largely abandoned in favor of an object-centered attentional state.

A note on asymmetry of color and motion attention effects

An examination of the scalp topographies and localizations across attention conditions in **Chapter 3** (**Figure 11**) raises the following question: why didn't we also observe the involvement of dorsal visual stream areas that could be attributed to enhancements of motion processing when motion (rather than color) was the attended feature? It might be the case that there is a genuine difference in the brain's capability to bias motion processing compared to color processing. This possibility was suggested by an earlier neuroimaging study by our group in which participants performed both a motion and a color task (Wylie et al., 2006). These two tasks were interlaced throughout each block such that some trials were repeats of the task on the prior trial whereas others required switching between tasks. Participants were slower to respond to motion targets when they had just done the color task (switch) than when they had just done the motion task (repeat), a so-called “switching cost.” Switching costs suggest that some time is needed to shift the focus of attention from one process (color) to another (motion). However, in this case the participants were given a minimum of 1.5 seconds of warning as to which task would be forthcoming, which is a fairly substantial interval in terms of the timescale of neural events. Presumably, participants would have the option to prepare the relevant neural circuitry in advance of the arrival of the target stimulus. Indeed, participants did not show switching costs when switching from the motion task to the *color* task, suggesting that they were able to effectively prepare color-related neural circuitry during the cue-target interval. This behavioral pattern was also reflected in the hemodynamic response: color-processing areas showed increased activity in preparation for the color task, but biasing activity in motion-processing areas was not detected in anticipation of the motion task. However, it may be the case that such differences in color- and motion-processing preparation are somewhat idiosyncratic. In **Chapter 1**, we did find group-level evidence for preparatory attention to both color and motion; however, we also found that nearly all participants selectively suppressed either color or motion, but not both. It could be that the clearer involvement of color processing areas in **Chapter 3** reflects the particular composition of the sample.

Idiosyncratic differences in feature-based biasing could provide an excellent fulcrum for further investigations of the anatomical basis of selective attention, and brain dynamics more generally. For example, in **Chapter 1**, we mentioned that participants' feature-selective oscillatory activities tended to covary with their idiosyncratic task performance biases. This admits the interesting possibility that the task performance biases are actually *caused by* a differential ability to engage various functionally-specialized cortical areas in particular dynamical states. In other words, some parts of cortex may not "resonate" as well with attentional control centers, which may result in less stable attentional control for that part of cortex and the cognitive processes it subserves. Variation in dynamical stability could be due to factors across a variety of scales ranging from single molecular species to global connectivity patterns. In the framework of this dynamical stability hypothesis, idiosyncratic differences could be exploited to infer which properties of brain organization are critical to supporting stable attentional control, insights that could be transferrable to other domains of brain function beyond attention.

Then again, it may not be the case that idiosyncratic oscillatory behavior *causes* the idiosyncracies of selective attention. An alternative is that the differences in task performance between features are due to some non-oscillatory cause, and the differences in alpha-mediated suppression reflect some sort of compensation. In other words, it could be the case that, due to some non-oscillatory cause, motion discrimination is more difficult than color discrimination for a given individual. Because motion discrimination is relatively difficult, the participants recruit more attentional resources for that task, reflected in greater electrophysiological measures of attentional biasing (i.e., alpha-band power increases). In this situation, although differences in brain dynamics are more of a consequence of idiosyncratic perceptual ability than a cause of idiosyncratic attentional deployment, such differences can still be used to identify subpopulations of “color suppressors” and “motion suppressors” for targeted inquiry into the bases of the more fundamental distinctions of brain operation between the two groups.

Future study: relationship of Gestalt grouping to binding

An important assumption underlying the experiments described in **Chapters 2** and **3** was that the STRONG object condition evokes greater binding processes. This was a working assumption and has not been explicitly demonstrated, although in retrospect the data support our contention. Nevertheless, we had *a priori* justification for believing that this Gestalt grouping manipulation would lead to differential engagement of binding. Key to our reasoning is the idea that an object must be *selected* by attention for binding to be effective (e.g., Treisman and Gelade, 1980). The principle argument against the idea that the STRONG and WEAK stimuli differentially invoke binding processes is that both stimuli have equally valid claims to being characterized as objects: one large object in the former case, and many smaller objects in the latter. The question then becomes that of whether one or more of the small objects comprising the WEAK stimulus can be attentionally selected, much as the one large object in the STRONG stimulus is selected. We believe this is unlikely for two reasons: (1) we designed our task such that attending to a single dot would not be an effective strategy, since identification of a target would require a comparison between a minimum of two dots, and (2) sequentially focusing attention on a series of dots to make such a comparison would also not be an effective strategy, since some time would be needed to processes one dot, disengage attention from that dot, shift attention to the next dot, and then process that dot (cf. Posner and Cohen, 1984). We presented the dot arrays for a 200 ms duration, which is the lower bound of covert attentional switch latencies by some estimates (e.g., Horowitz, Holcombe, Wolfe, Arsenio and DiMase, 2004). Even if two spatial shifts of attention were possible during the stimulus duration, successful target identification would require picking two dots that differed from the target feature, which was unlikely (about 16%). Because of this, we believe that participants most likely directed attention to features, but not objects, in order to attain adequate performance for WEAK stimuli. Despite these reasons why we believe that an object-centered attentional state is supported for STRONG stimuli but not for WEAK stimuli, future studies could be aimed at directly investigating if our Gestalt grouping maniulation modulates the robustness of binding processes.

One challenge to demonstrating directly that the STRONG object leads to more robust binding processes is that binding remains a fairly subjective phenomenon. Whereas modern neuroimaging and electrophysiological methods can assay to some extent which features an experimental subject is viewing (and even where in the visual field those features are perceived) this alone is not sufficent to establish binding, which requires the assignment of various features to a common *object*. This latter aspect of binding has not been definitively linked to an objective electrophysiological indicator. Because a definitive neurophysiological indicator of binding is not yet known, direct evidence of binding has so far come from patterns of behavior. In particular, certain patterns of errors in visual search tasks are consistent with participants forming “illusory conjunctions” of features that actually belong to different objects (Treisman and Gelade, 1980). Under different experimental conditions the incidence of such illusory conjunctions can be modulated, indicating that the robustness of binding processes is dependent on the experimental manipulation. For example, Treisman and Schmidt (1982) showed evidence that stressing spatial attention demands increases the incidence of illusory conjunctions, which led these researchers to conclude that successful binding is a spatial attention-dependent process.

One future study that could be directed at investigating directly the role of our Gestalt grouping manipulation on binding processes is a variant of the visual search paradigms used by Anne Treisman and colleagues (e.g., Treisman and Gelade, 1980; Treisman and Schmidt, 1982). The search array in the proposed experiment would be populated with apertures containing random dot stimuli like those described in this report, albeit smaller. If our assumption that STRONG stimuli invoke more robust binding than do WEAK stimuli is true, then we would predict that illusory conjunctions would be more prevalent when the search array consists of WEAK objects than when it consists of STRONG objects.

Future study: role of suppression in feature-based gain and tuning

In **Chapter 1**, we reported differential modulation of alpha-band activity when attending to different features. For example, alpha power from dorsal visual stream regions increased more when color was attended than when motion was attended. We interpreted this pattern of results as consistent with the known role of alpha band activity as an active suppression mechanism in spatial selective attention. One interesting observation, however, is that alpha power increases during the preparatory interval regardless of which feature is attended. With respect to the earlier example, alpha power from dorsal visual stream regions increased more when color was attended than when motion was attended, *but it increased relative to baseline in both cases*. It may seem that this latter pattern of results is inconsistent with the presumptive role of alpha-band increases as a suppressive mechanism. However, as described below, it may be the case that some suppression can actually *improve* feature discrimination ability. Future research can be directed to test this hypothesis.

To interpret the “non-specific” alpha power increases described above, it is useful to first have some notion of what, exactly, might be suppressed in this context. One popular model of attentional processes is the so-called “biased competition” model. In this model, converging information streams are faced with a capacity limitation (a “bottleneck”), and attention biases the “competition” between the information streams such that behaviorally relevant information gets processing priority. Attention can achieve this biasing by adjusting the quantity or quality of competing information. The former process is “gain modulation” (quantity) and the latter is “tuning” (quality). This distinction is important because the role of suppression has different implications for the two modes of biasing. In the gain modulation mode, suppression reduces the quantity of information, whereas in the tuning mode, suppression can *improve* the quality of information. **Figure 13** illustrates this point. Each curve in the figure represents the response function of a hypothetical direction-selective neuron or neural assembly. The dashed curves represent the response function prior to attentional modulation, and the solid curves are the biased responses. The region bounded in red represents the portion of the response curve that would be suppressed for the biasing mode indicated. In the gain case the entire process is suppressed, whereas only portions of the process are suppressed in the tuning case. This response curve narrowing in the tuning case translates to an improved ability of the process to discriminate the “ideal” direction from others. This has important implications for the role of suppression in feature-based attention.

Whereas spatial attention is consistently shown to involve gain modulation without response selectivity tuning, both gain and tuning are implicated in feature-based attention (Ling, Liu and Carrasco, 2009; Martinez-Trujillo and Treue, 2004). This reflects the essentially discriminatory nature of feature-based orientations (e.g., discriminating a relevant color from an irrelevant one). The following example emphasizes the role of suppression in attention-mediated tuning. Martinez-Trujillo and Treue (2004) trained monkeys to attend to patterns of moving dots while they recorded from neurons with receptive fields (RFs) outside the focus of the animals’ attention. The receptive fields of the recorded neurons also contained moving dot stimuli, and these stimuli could move in the same direction as the attended stimuli, or a different direction. They found that attending to the moving dots enhanced responses to RF stimuli moving in a similar direction, but *suppressed* responses to stimuli moving in directions that were quite different from that of the attended stimulus. In other words: although motion was attended (monkeys were responding to changes in direction), at least some suppression was found for cells processing motion. Several mechanisms for this suppression are conceivable, but if alpha-band oscillations play any role, then it would not be surprising to observe alpha-band power increases over cortical areas processing the to-be-attended feature modality (motion processing areas in the dorsal visual stream, for this example).

Thus, there is the potential that suppression acts on both the cortical areas processing the relevant feature (tuning) and those areas processing the irrelevant feature (gain). This potentially complicates the interpretation of the results from **Chapter 1** in isolation. A future study could aim to disambiguate the potential contributions of gain and tuning to the observations at the scalp. One way to do this could be to lower the exposure duration of the stimulus while maintaining performance at 80% through our titration procedure. The discrimination threshold for motion direction (Conklin, Baldwin and Brown, 1953; Festa and Wesch, 1997) and hue (Hita, Romero, Jimenez del Barco and Martinez, 1982; Siegel, 1965) are both inversely related to stimulus duration, and so the two values for colors and motion directions in the targets will be more dissimilar than with the longer exposure duration in **Chapter 1** (**Figure 14**). Because the feature values will be less similar, there will be less competition for processing resources within a feature area (e.g., less competition between processing of particular color values). However, since performance is held constant, a similar degree of capacity limitation is presumably maintained, and there should be a commensurate increase in competition between feature modalities (e.g., more competition between processing of color and processing of motion), which is consistent with the gain model.

By taking two measurements of alpha-band suppression effects with different demands on tuning and gain modulation modes of biasing, we will be able to provide a richer interpretation of the roles of these two processes in feature-based attention than would be possible with one experiment alone. If, for example, anticipatory increases in brain regions processing the to-be-attended feature are attenuated when stimulus durations are shortened, then this would suggest that such increases are at least partially related to tuning control.

On "objectness": a competing cluster hypothesis

A central component to this thesis was an attempt to manipulate the "objectness" of visual events. We might step back for a moment and ask what exactly is meant by "objectness" in this context, and what it would mean to have different degrees of it. The term "objectness" here does not refer to things in the external environment, such as those we call chairs, gophers or tornados. Rather, "objectness" refers to a quality of the brain's *representation* of the things in the external environment (i.e., the *percept* of a tornado). Indeed, the notion of a material "object" does not exist without someone to *perceive* that object. Ironically, there can be no *objective* definition of an *object*. Most operational definitions of an object are variations on this concise version from philosopher Charles Sanders Peirce (1839-1914): "By an object, I mean anything that we can think, i.e. anything we can talk about" (n.d.). Something becomes an object by virtue of someone perceiving or thinking about it as such: by asserting a categorical distinction between the properties that the perceiver assigns to that object from other properties of the environment. These asserted boundaries are arbitrary and artificial: there are no categorical boundaries *per se* between a tornado (or a chair, for that matter) and the rest of the atmosphere.

Thus, what we aim to manipulate when we speak of manipulating "objectness" is a quality of a person's *perception* of a stimulus. When a stimulus is perceived as an object, its features are conceptually bound together and segregated from other features in the environment. When a stimulus is not perceived as an object, this sort of categorical assignment of features does not occur. Presumably there is room for continuity between these two extremes reflected by intermediate degrees of categorization. A key question of contemporary brain science is *what are the physical mechanisms that bind features together* (or the oft-neglected corollary: *what are the mechanisms that keep features segregated*)?

Of course, terms such as “bind together” are metaphorical here. There are very few features of a tornado that are physically present in a brain, besides it being grey and wet. Rather, a brain perceiving a tornado is characterized by patterns of activity that *represent* the features of a tornado. The concept of binding also refers to a pattern of activity, but one that represents an *association* between the features of the object. Given that the representations of the separate features themselves typically rely on relatively distinct functionally specialized cortical areas, such a representation of association likely involves coordinated activity across these areas.

One way to coordinate activity across cortical areas is by convergent inputs to a common computational unit (e.g., the so-called “grandmother cell”), the activity of which could represent the association between the features. However, this solution is unlikely to be sufficiently flexible to represent the near infinite variety of objects in nearly infinite configurations encountered in the course of a typical lifetime –especially since a bump on the head could erase all ability to conceive of your grandmother unless substantial redundancy was involved. Moreover, the question remains as to *how* the activity of that one computational unit comes to *represent* anything to the animal. A reasonable argument could be that the activity of the cell influences the activity of other cells that in turn contribute to behavior, such that the animal behaves as though it perceives the object in question. However, this interpretation belies the “grandmother cell” concept, as it invokes a network-level description involving coordination between multiple units – which brings us to the distributed representation account of binding. This account posits that the pattern of activity that represents the association between features of an object is not confined to a single computational unit, but rather distributed across a large-scale network.

In this section we have been trying to operationalize what is meant by “objectness”. So far we have argued that objectness is subjectively-generated and not naturally-occurring; that it involves the association of some features and the segregation of others; and that it is a pattern of activity across a large-scale cortical network. The next logical question is *what is the nature of this coordinated, distributed activity that represents the association between individual features*? As mentioned in the previous chapters, one promising candidate is oscillatory synchrony between brain areas. However, for the purposes of this discussion it is not necessary to endorse a particular class of dynamics: the “binding” activity can be conceived as periodic, chaotic (deterministically so, of course), or some intermediate degree of complexity. Some general conclusions about brain operation can be made with incomplete knowledge of the dynamics involved.

First of all, it is clear that only a small number of binding “activities” can stably exist in one brain at one time. While our perception gives the illusion of an experience completely populated with discrete objects, in fact we have conscious *access* to only a handful of objects at a time. This claim is supported by results from studies of multiple object tracking, illusory feature conjunctions, working memory span, change blindness, subitizing (rapid, non-counting judgment of number), spatial reorientation, divided attention and many other phenomena. By “conscious access”, it is meant simply that the person can think about the object, or in other words, they can entertain a proposition about that object. In fact, if the person can’t think about something it is *ipso facto* disqualified from objecthood (see Peirce’s definition above).

This limitation on the number of objects that can be perceived at a time can be counterintuitive because people do not generally report having the experience that objects become “unbound” the moment they cease to think about them. The catch, however, is that the very act of introspecting “did that X become unbound?” ensures that the X snaps back together as a fully-formed object (by definition). Further complicating matters is the brain’s remarkable ability to confabulate, such that large swaths of our experience can be filled in with a mnemonic brush to create an illusion of stability. Nevertheless, careful experiments, and popular visual demonstrations[1], routinely show that we can truly perceive only a few objects at a time.

What does this limit on the number of objects that can be simultaneously perceived reveal about the basis of brain function? We have already introduced the concept that binding involves distributed activity across a large-scale network. The set of dynamics that can be supported by a network is dependent on the shape, or topology of that network. Thus, the limitations on supported dynamics can be used to infer properties of the network topology. The topology of a network is the set of pairwise connections (“links” in the terminology of graph theory) between the computational units (“nodes”). Even without a complete fine-grained description of the topology, important properties of network dynamics can be inferred from coarse-grained statistical descriptions of the *distribution* of links in the network. For example, one particularly useful metric for

---

[1] e.g., http://www.psychology.uoguelph.ca/faculty/trick/motCSS.html

describing a network is the clustering coefficient, which quantifies how densely interconnected is a network or a part of a network. A network with a high global clustering coefficient is unlikely to support multiple independent activity patterns in different subsets of the network by virtue of their high interconnectedness. On the other hand, any stable dynamics (e.g., binding activity) across a subset of the network is unlikely to be supported without *some* local clustering to stabilize it ("local" refers to the particular subset of the network, which can itself be widely-distributed in spatial terms). That is, if a given subpart of the network is as connected to other parts of the network as it is to itself, its dynamics are likely to be fairly susceptible to external perturbations.

We suggest that binding activity can be conceived of as a relatively stable dynamic activity pattern in a cluster, likely spanning feature-specified cortical areas. A few such clusters can be supported at a time with some degree of independence in their activities. However, the clusters are not fully independent, as they remain connected to each other and the rest of the brain. Moreover, the clustering must be dynamically modifiable to deal with the near infinite number of objects that can be perceived without positing an impossibly large number of fixed clusters (the very problem with the "grandmother cell" model). Perhaps as the number of independent binding activities increases, the chance of destabilizing cross-talk between the clusters via their "loose" inter-cluster connections also increases. This might account for why only a few objects can be perceived at a time.

A "competing cluster" hypothesis regarding object perception would predict that selective attention should be more effective in biasing processing between two or more such clusters than between the constituent features *within* one such cluster. This is because the cluster is defined by its high interconnectedness, and connectedness in network theoretical terms is not merely a structural concept. Rather, a connection is only meaningful to network analysis if it is *influential* to the behavior of the connected nodes. Thus, the state of each node of a cluster is highly dependent on the states of all the other nodes, encouraging a certain amount of parity of processing within the cluster. On the other hand, two separate clusters are separate precisely because of the relative independence of their activities. This means that selective attention should readily ramp up or down the quality of processing across an entire cluster[2], but should be less effective in modulating processing of only a part of a cluster. This prediction was precisely what was tested, and confirmed, by this thesis.

A model of feature-based attentional deployment

The evidence from **Chapter 1** suggests that anticipatory attention can be deployed with respect to basic visual features much like it is deployed to spatial locations. This biased preparatory state is set up prior to the arrival of the stimulus. As such, it could be an effective way to “key in” on a relevant feature for the purposes of guiding visual search (Egner et al., 2008; Hopf et al., 2004). However, the evidence from **Chapters 2** and **3** suggest that the attentional state of the brain undergoes a rapid and dramatic reconfiguration once an object comes into view.

---

[2] Note that a natural way to adjust the quality of processing for an object in this framework is to stabilize or destabilize the cluster. This could be accomplished by changing the strength of a few connections, which can have widespread effects on clustering throughout the network.

When the stimulus in our experiments is not readily perceived as a single coherent object, differences in stimulus-processing are evident at relatively early timeframes (~180 ms). These differences are attributable to enhanced activity in color-related regions when color is attended compared to when motion is attended. This indicates that feature-based anticipatory attention can result in behaviorally meaningful stimulus processing differences –at least under certain conditions; specifically when attention is not directed to an *object*.

When the stimulus in our experiments can be perceived as a single coherent object, however, both the behavioral advantages of cue information (**Chapter 2**) and electrophysiological correlates of biased perceptual processing (**Chapter 3**) disappear. What this suggests is that once an object (rather than a mere feature) is selected by attention, the biased processing state is abandoned in favor of object-related binding processes, which seem to involve a spread of attention across the constituent features of the object (Albrecht et al., 2008; Egly et al., 1994; Fiebelkorn et al., 2010a, 2010b; Katzner et al., 2009; Molholm et al., 2007; O'Craven et al., 1999; Schoenfeld et al., 2003; Wylie et al., 2004). That is, once an object is localized, the feature-based approach is replaced by an object-based one. This shift in the attentional "reference frame" happens prior to 180 ms. The attended object can be scrutinized with respect to a particular feature once the binding is established, but this is a different process than the early perceptual effects due to preparatory feature-based biasing (**Chapter 3**).

Figure 13 (caption): Models of suppression in gain and tuning response modulations. The dashed curves represent the response function prior to attentional modulation, and the solid curves are the biased responses. The region bounded in red represents the portion of the response curve that would be suppressed for the biasing mode indicated.

Figure 13:

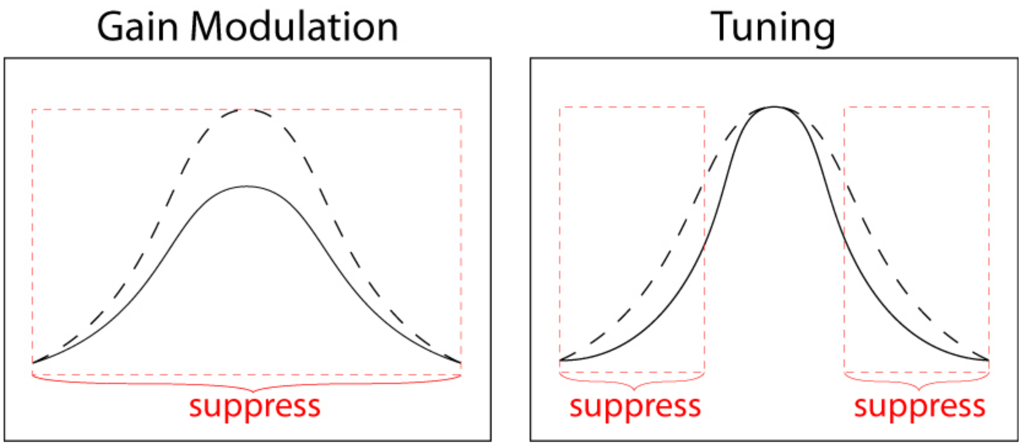

Figure 14 (caption): Schematized illustration of discrimination/detection load trade-off. Each curve is the function of performance versus exposure duration for a discrimination task of a given difference. By specifying the duration and performance and leaving the difference free to vary, the desired parameters will be attained through titration.

Figure 14:

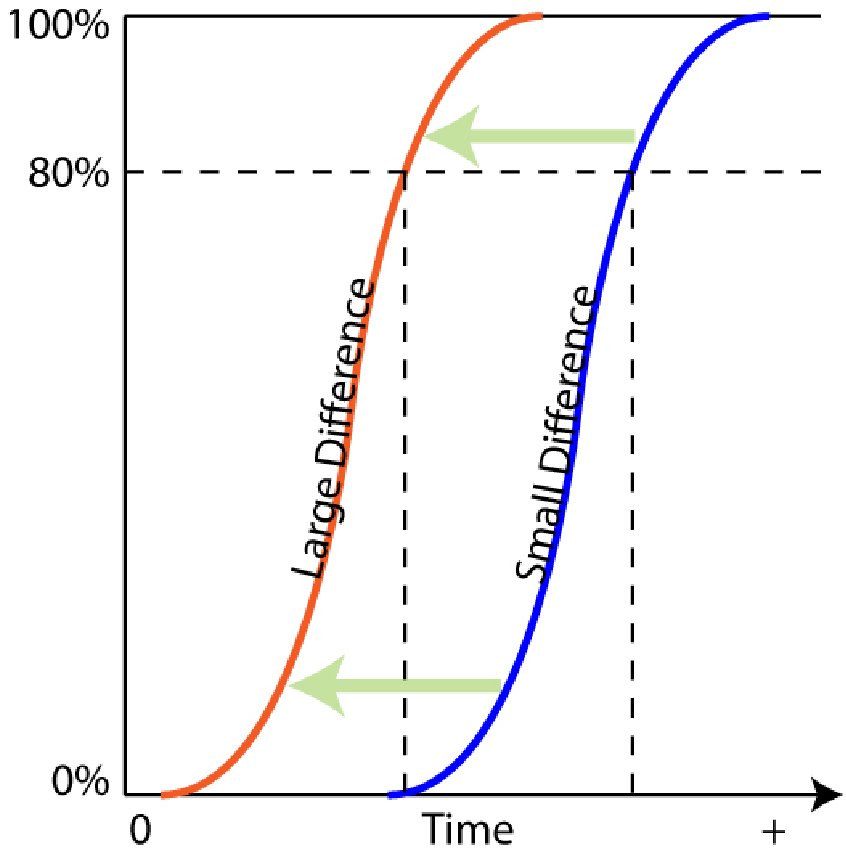

REFERENCES

Ahlsén, G., Lindström, S., & Lo, F. S. (1985). Interaction between inhibitory pathways to principal cells in the lateral geniculate nucleus of the cat. *Experimental Brain Research, 58*(1), 134-143.

Albrecht, A. R., List, A., & Robertson, L. C (2008). Attentional selection and the representation of holes and objects. *Journal of Vision, 8*(13), 1-10.

Anllo-Vento, L., & Hillyard, S. A. (1996). Selective attention to the color and direction of moving stimuli: electrophysiological correlates of hierarchical feature selection. *Perception & Psychophysics, 58*(2), 191-206.

Armstrong, C. L., Hayes, K. M., & Martin, R. (2001). Neurocognitive Problems in Attention Deficit Disorder. *Annals of the New York Academy of Sciences, 931*, 196–215. doi: 10.1111/j.1749-6632.2001.tb05780.x

Banerjee, S., Snyder, A. C., Molholm, S., & Foxe, J. J. (2011). Oscillatory alpha-band mechanisms and the deployment of spatial attention to anticipated auditory and visual target locations: supramodal or sensory-specific control mechanisms? *Journal of Neuroscience, 31*(27), 9923-9932.

Berger, H. (1929). Über das Elektroenkephalogramm des Menschen. *Archiv für Psychiatrie und Nervenkrankheiten, 87,* 527-570.

Blaser, E., Pylyshyn, Z. W., & Holcombe , A. O. (2000). Tracking an object through feature space. *Nature*, *408*, 196-199.

Bollimunta, A., Chen, Y., Schroeder, C. E., & Ding, M. (2008). Neuronal mechanisms of cortical alpha oscillations in awake-behaving macaques. *Journal of Neuroscience, 28*(40), 9976-9988.

Brandt, M. E. (1997). Visual and auditory evoked phase resetting of the alpha EEG. *International Journal of Psychophysiology, 26*(1-3), 285-298.

Brandt, M. E., & Jansen, B. H. (1991). The relationship between prestimulus-alpha amplitude and visual evoked potential amplitude. *International Journal of Neuroscience, 61*(3-4), 261-268.

Brandt, M. E., Jansen, B. H., & Carbonari, J. P. (1991). Pre-stimulus spectral EEG patterns and the visual evoked response. *Electroencephalography and Clinical Neurophysiology*, *80*(1), 16-20.

Broadbent, D. (1958). *Perception and Communication.* London, UK: Pergamon Press.

Busse, L., Roberts, K. C., Crist, R. E., Weissman, D. H., & Woldorff, M. G. (2010). The spread of attention across modalities and space in a multisensory object. *Proceedings of the National Academy of Sciences USA, 102*(51), 18751-18756.

Capotosto, P., Babiloni, C., Romani, G. L., & Corbetta, M. (2009). Frontoparietal cortex controls spatial attention through modulation of anticipatory alpha rhythms. *Journal of Neuroscience, 29*(18), 5863-5872.

Conklin, J. E., Baldwin, A., & Brown, R. H. (1952). Apparatus for measuring the threshold for visual discrimination of direction of movement. *The American Journal of Psychology, 66*(2), 289-294.

Contreras, D., Destexhe, A., Sejnowski, T. J., & Steriade M. (1996). Control of spatiotemporal coherence of a thalamic oscillation by corticothalamic feedback. *Science, 274*(5288), 771-774.

Corbetta, M., Miezin, F. M., Dobmeyer, S., Shulman, G. L., & Petersen, S. E. (1991). Selective and divided attention during visual discriminations of shape, color, and speed: functional anatomy by positron emission tomography. *Journal of Neuroscience, 11*(8), 2383-2402.

Corbetta, M., Miezin, F. M., Shulman, G. L., & Petersen, S. E. (1993). A PET study of visuospatial attention. *Journal of Neuroscience, 13*(3), 1202-1226.

Corbetta, M., Kincade, J. M., Ollinger, J. M., McAvoy, M. P., & Shulman, G. L. (2000). Voluntary orienting is dissociated from target detection in human posterior parietal cortex. *Nature Neuroscience, 3*(3), 292-297.

Cosmelli, D., López, V., Lachaux, J.-P., López-Calderón, J., Renault, B., Martinerie, J., & Aboitiz, F. (2011). Shifting visual attention away from fixation is specifically associated with alpha band activity over ipsilateral parietal regions. *Psychophysiology, 48*(3), 312-322.

Dalton, P., & Lavie, N. (2004). Auditory attentional capture: effects of singleton distractor sounds. *Journal of Experimental Psychology: Human Perception and Performance, 30*(1), 180-193.

Delorme, A., Sejnowski, T., & Makeig, S. (2007). Enhanced detection of artifacts in EEG data using higher-order statistics and independent component analysis. *NeuroImage, 34*, 1443-1449.

Derrington, A. M., Krauskopf, J., & Lennie, P. (1984). Chromatic mechanisms in lateral geniculate nucleus of macaque. *Journal of Physiology, 357*, 241-265.

Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience, 18,* 193-222.

Desimone, R., & Ungerleider, L. (1986). Multiple visual areas in the caudal superior temporal sulcus of the macaque. *Journal of Comparative Neurology, 248,* 164-189.

Destexhe, A., Contreras, D., Steriade, M. (1998). Mechanisms underlying the synchronizing action of corticothalamic feedback through inhibition of thalamic relay cells. *Journal of Neurophysiology, 79*(2), 999-1016.

Dockree, P. M., Kelly, S. P., Foxe, J. J., Reilly, R. B., & Robertson I. H. (2007). Optimal sustained attention is linked to the spectral content of background EEG activity: greater ongoing tonic alpha (approximately 10 Hz) power supports successful phasic goal activation. *European Journal of Neuroscience, 25*(3), 900-907.

Doniger, G. M., Foxe, J. J., Murray, M. M., Higgins, B. A., Snodgrass, J. G., Schroeder, C. E., & Javitt, D. (2000). Activation timecourse of ventral visual stream object-recognition areas: high density electrical mapping of perceptual closure processes. *Journal of Cognitive Neuroscience, 12*(4), 615-621.

Doniger, G. M., Foxe, J. J., Schroeder, C. E., Murray, M. M., Higgins, B. A., & Javitt, D. C. (2001). Visual perceptual learning in human object recognition areas: a repetition priming study using high-density electrical mapping. *Neuroimage, 13*(2), 305-313.

Dufour, A. (1999). Importance of attentional mechanisms in audiovisual links. *Experimental Brain Research, 126*(2), 215-222.

Eckhorn, R., Bauer, R., Jordan, W., Brosch, M., Kruse, W., Munk, M., & Reitboeck, H. J. (1988). Coherent oscillations: a mechanism of feature linking in the visual cortex? Multiple electrode and correlation analyses in the cat. *Biological Cybernetics, 60*(2), 121-130.

Eglin, M., Robertson, L. C., & Knight, R. T. (1989). Visual search performance in the neglect syndrome. *Journal of Cognitive Neuroscience, 4,* 372-381.

Eglin, M., Robertson, L. C., & Knight, R. T. (1991). Cortical substrates supporting visual search in humans. *Cerebral Cortex, 1,* 262-272.

Eglin, M., Robertson, L. C., Knight, R. T., & Brugger, P. (1994). Search deficits in neglect patients are dependent on the size of the visual scene. *Neuropsychology, 4,* 451-463.

Egly, R., Driver, J., & Rafal, R. D. (1994). Shifting visual attention between objects and locations: evidence from normal and parietal lesion subjects. *Journal of Experimental Psychology, 123*(2), 161-177.

Egner, T., Monti, J. M. P., Trittschuh, E. H., Wieneke, C. A., Hirsch, J., & Mesulam, M.-M. (2008). Neural integration of top-down spatial and feature-based information in visual search. *Journal of Neuroscience, 28*(24), 6141-6151.

Eimer, M. (1999). Can attention be directed to opposite locations in different modalities? An ERP study. *Clinical Neurophysiology, 110*(7), 1252-1259.

Eimer, M., & Schröger, E. (1998). ERP effects of intermodal attention and cross-modal links in spatial attention. *Psychophysiology, 35*(3), 313-327.

Esterman, M., McGlinchey-Berroth, R., & Milberg, W. P. (2000). Parallel and serial search in hemispatial neglect: evidence for preserved preattentive but impaired attentive processing. *Neuropsychology, 14,* 599-611.

Farah, M. J., Wong, A. B., Monheit, M. A., Morrow, L. A. (1989). Parietal lobe mechanisms of spatial attention: modality-specific or supramodal? *Neuropsychologia, 27*(4), 461-470.

Felleman, D. J., & Van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex, 1*(1), 1-47.

Festa, E. K., & Wesch, L. (1997). Recruitment mechanisms in speed and fine-direction discrimination tasks. *Vision Research, 37*(22), 3129-3143.

Fiebelkorn, I. C., Foxe, J. J., & Molholm, S. (2010a). Dual mechanisms for the cross-sensory spread of attention: how much do learned associations matter? *Cerebral Cortex*, *20*(1), 109-120.

Fiebelkorn, I. C., Foxe, J. J., Schwartz, T. H., & Molholm, S. (2010b). Staying within the lines: the formation of visuospatial boundaries influences multisensory feature integration. *European Journal of Neuroscience, 31*(10), 1737-1743.

Foxe, J. J., Simpson, G. V., & Ahlfors, S. P. (1998). Parieto-occipital ~10 Hz activity reflects anticipatory state of visual attention mechanisms. *Neuroreport, 9*, 3929-3933.

Foxe, J. J., McCourt, M. E., & Javitt, D. C. (2003). Right hemisphere control of visuospatial attention: line-bisection judgments evaluated with high-density electrical mapping and source analysis. *NeuroImage, 19*(3), 710-726.

Foxe, J. J., Simpson, G. V., Ahlfors, S. P., & Saron, CD. (2005a), Biasing the brain's attentional set: I. Cue driven deployments of intersensory selective attention. *Experimental Brain Research, 166*, 370-392.

Foxe, J. J., Murray, M. M., & Javitt, D. C. (2005b). Filling-in in schizophrenia: a high-density electrical mapping study and source-analysis investigation of illusory contour processing. *Cerebral Cortex, 15*(12), 1914-1927.

Fries, P., Reynolds, J. H., Rorie, A. E., & Desimone, R. (2001). Modulation of oscillatory neuronal synchronization by selective visual attention. *Science, 291*(5508), 1560-1563.

Fu, K. M., Foxe, J. J., Murray, M. M., Higgins, B. A., Javitt, D. C, & Schroeder, C. E. (2001). Attention-dependent suppression of distracter visual input can be cross-modally cued as indexed by anticipatory parieto-occipital alpha-band oscillations. *Brain Research: Cognitive Brain Research, 12*(1), 145-152.

Gomez-Ramirez, M., Higgins, B. A., Rycroft, J. A., Owen, G. N., Mahoney, J., Shpaner, M., & Foxe, J. J. (2007). The deployment of intersensory selective attention: a high-density electrical mapping study of the effects of theanine. *Clinical Neuropharmacology, 30*(1), 25-38.

Gomez-Ramirez, M., Kelly, S. P., Montesi, J. L., & Foxe, J. J. (2009). The effects of L-theanine on alpha-band oscillatory activity during a visuo-spatial attention task. *Brain Topography, 22*(1), 44-51.

Gould, I. C., Rushworth, M. F., & Nobre, A. C. (2011). Indexing the graded allocation of visuospatial attention using anticipatory alpha oscillations. *Journal of Neurophysiology,* Epub ahead of print. doi: 10.1152/jn.00653.2010

Graimann, B., & Pfurtscheller, G. (2006). Quantification and visualization of event-related changes in oscillatory brain activity in the time-frequency domain. In C. Neuper & W. Klimesch (Eds.), *Progress in Brain Research* (79-97). Amsterdam: Elsevier.

Gray, C. M., König, P., Engel, A. K., & Singer, W. (1989). Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature, 338*, 334-337.

Gruber, T., Müller, M. M., Keil, A., & Elbert, T. (1999). Selective visual-spatial attention alters induced gamma band responses in the human EEG. *Clinical Neurophysiology, 110*(12), 2074-2085.

Guthrie, D., & Buchwald, J. S. (1991). Significance testing of difference potentials. *Psychophysiology*, *28*(2), 240-244.

Haenny, P. E., & Schiller, P. H. (1988). State dependent activity in monkey visual cortex. *Experimental Brain Research*, *69*(2), 225-244.

Hanslmayr, S., Aslan, A., Staudigl, T., Klimesch, W., Herrmann, C. S., & Bäuml, K.-H. (2007). Prestimulus oscillations predict visual perception performance between and within subjects. *NeuroImage, 37*(4), 1465-1473.

Heilman, K. M., & Van Den Abell, T. (1980). Right hemisphere dominance for attention: the mechanism underlying hemispheric asymmetries of inattention (neglect). *Neurology, 30*(3), 327-330.

Herrmann, C. S., Mecklinger, A., & Pfeifer, E. (1999). Gamma responses and ERPs in a visual classification task. *Electroencephalography and Clinical Neurophysiology, 110*, 636-642.

Hillyard, S. A., & Anllo-Vento, L. (1998). Event-related brain potentials in the study of visual selective attention. *Proceedings of the National Academy of Sciences USA, 95*(3), 781-787.

Hillyard, S. A., Vogel, E. K., & Luck, S. J. (1998). Sensory gain control (amplification) as a mechanism of selective attention: electrophysiological and neuroimaging evidence. *Philosophical Transactions of the Royal Society of London B Biological Sciences, 353*(1373), 1257-1270.

Hita, E., Romero, J., Jimenez del Barco, L., & Martinez, R. (1982). Temporal aspects of color discrimination. *Journal of the Optical Society of America, 72*(5), 578-582.

Hopf, J.-M., Boelmans, K., Schoenfeld, M. A., Luck, S. J., & Heinze, H.-J. (2004). Attention to features precedes attention to locations in visual search: evidence from electromagnetic brain responses in humans. *Journal of Neuroscience, 24*(8), 1822-1832.

Horowitz, T. S., Holcombe, A. O., Wolfe, J. M., Arsenio, H. C., & DiMase, J. S. (2004). Attentional pursuit is faster than attentional saccade. *Journal of Vision, 4,* 585-603.

Hyvärinen, A. (1999). Fast and robust fixed-point algorithms for independent component analysis. *IEEE Transactions on Neural Networks, 10*(3), 626-634.

Ivanoff, J., & Saoud, W. (2009). Nonattentional effects of nonpredictive central cues. *Attention, Perception, & Psychophysics, 71,* 872-880.

Jacobson, S., & Trojanowski, J. Q. (1975). Corticothalamic neurons and thalamocortical terminal fields: an investigation in rat using horseradish peroxidase and autoradiography. *Brain Research, 85*(3), 385-401.

Jahnsen, H., & Llinás, R. (1984). Ionic basis for the electro-responsiveness and oscillatory properties of guinea-pig thalamic neurons in vitro. *Journal of Physiology, 349,* 227-247.

Jokisch, D., & Jensen, O. (2007). Modulation of gamma and alpha activity during a working memory task engaging the dorsal or ventral stream. *Journal of Neuroscience, 27*(12), 3244-3251.

Jones, S. R., Kerr, C. E., Wan, Q., Pritchett, D. L., Hämäläinen, M., & Moore, C. I. (2010). Cued spatial attention drives functionally relevant modulation of the mu rhythm in primary somatosensory cortex. *Journal of Neuroscience, 30*(41), 13760-13765.

Kahneman, D. (1973). *Attention and Effort.* Englewood Cliffs, NJ: Prentice-Hall.

Kanizsa, G. (1976). Subjective Contours. *Scientific American, 234,* 48-52.

Katzner, S. Busse, L., & Treue, S. (2009). Attention to the color of a moving stimulus modulates motion-signal processing in macaque area MT: evidence for a unified attentional system. *Frontiers in Neuroscience, 3*(12), 1-8.

Keil, A., & Müller, M. M. (2010). Feature selection in the human brain: electrophysiological correlates of sensory enhancement and feature integration. *Brain Research, 1313,* 172-184.

Kelly, S. P., Lalor, E. C, Reilly, R. B., & Foxe, J. J. (2005). Visual spatial attention tracking using high-density SSVEP data for independent brain-computer communication. *IEEE Transactions on Neural Systems and Rehabilitation Engineering, 13*(2), 172-178.

Kelly, S. P., Lalor, E.C, Reilly, R. B., & Foxe, J. J. (2006). Increases in alpha oscillatory power reflect an active retinotopic mechanism for distracter suppression during sustained visuospatial attention. *Journal of Neurophysiology, 95*(6), 3844-3851.

Kelly, S. P., Gomez-Ramirez, M., & Foxe, J. J. (2008). Spatial attention modulates initial afferent activity in human primary visual cortex. *Cerebral Cortex, 18*(11), 2629-2636.

Kelly, S. P., Gomez-Ramirez, M., & Foxe, J. J. (2009). The strength of anticipatory spatial biasing predicts target discrimination at attended locations: a high-density electroencephalographic study. *European Journal of Neuroscience*, *30*, 2224-2234.

Kelly, S. P., Foxe, J. J., Newman, G., & Edelman, J. A. (2010). Prepare for conflict: EEG correlates of the anticipation of target competition during overt and covert shifts of visual attention. *European Journal of Neuroscience, 31*(9), 1690-1700.

Kerlin, J. R., Shahin, A. J., & Miller, L. M. (2010). Attentional gain control of ongoing cortical speech representations in a “cocktail party”. *Journal of Neuroscience, 30*(2), 620-628.

Klimesch, W., Doppelmayr, M., Russegger, H., Pachinger, T., & Schwaiger, J. (1998). Induced alpha band power changes in the human EEG and attention. *Neuroscience Letter, 244*(2), 73-76.

Koch, K., McLean, J., Segev, R., Freed, M. A., Berry II, M. J., Balasubramanian, V., & Sterling, P. (2006). How much the eye tells the brain. *Current Biology, 16,* 1428-1434.

LaBerge, D. (1997). Attention, awareness, and the triangular circuit. *Consciousness and Cognition, 6*(2-3), 149-181.

Laeng, B., Brennen, T., Espeseth, T. (2002). Fast responses to neglected targets in visual search reflect pre-attentive processes: an exploration of response times in visual neglect. *Neuropsychologia, 40,* 1622-1636.

Lindgren, K. A., Larson, C. L., Schaefer, S. M., Abercrombie, H. C, Ward, R. T., Oakes, T. R., Holden, J. E., Perlman, S. B., Benca, R. M., & Davidson, R. J. (1999). Thalamic metabolic rate predicts EEG alpha power in healthy control subjects but not in depressed patients. *Biological Psychiatry, 45*(8), 943-952.

Ling, S., Liu, T., & Carrasco, M. (2009). How spatial and feature-based attention affect the gain and tuning of population responses. *Vision Research, 49,* 1194-1204.

Lipp, O. V., & Waters, A. M. (2007). When danger lurks in the background: attentional capture by animal fear-relevant distracters is specific and selectively enhanced by animal fear. *Emotion, 7*(1), 192-200.

Liu, T., Slotnick, S. D., Serences, J. T., & Yantis, S. (2003). Cortical mechanisms of feature-based attentional control. *Cerebral Cortex, 13,* 1334-1343.

Livingstone, M. & Hubel, D. (1988). Segregation of form, color, movement & depth: anatomy, physiology and perception. *Science, 275*, 671-674.

Lopes da Silva, F. (1991). Neural mechanisms underlying brain waves: from neural membranes to networks. *Electroencephalography and Clinical Neurophysiology, 79,* 81-93.

Lopes da Silva, F. H., Vos, J. E., Mooibroek, J., & Van Rotterdam, A. (1980). Relative contributions of intracortical and thalamo-cortical processes in the generation of alpha rhythms, revealed by partial coherence analysis. *Electroencephalography and Clinical Neurophysiology*, *50*(5-6), 449-456.

Lucan, J. N., Foxe, J. J., Gomez-Ramirez, M., Satian, K., & Molholm, S. (2010). Tactile shape discrimination recruits human lateral occipital complex during early perceptual processing. *Human Brain Mapping, 31*(11), 1813-1821.

Luck, S. J., Chelazzi, L., Hillyard, S. A., & Desimone, R. (1997). Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. *Journal of Neurophysiology, 77*(1), 24-42.

Markram, K., & Markram, H. (2010). The intense world theory – a unifying theory of the neurobiology of autism. *Frontiers in Human Neuroscience,4,* 204. Doi: 10.3389/fnhum.2010.00224.

Martínez, A., Ramanathan, D. S., Foxe, J. J., Javitt, D. C., & Hillyard, S. A. (2007b). The role of spatial attention in the selection of real and illusory objects. *Journal of Neuroscience, 27*(3), 7963-7973.

Martínez, A., Teder-Sälejärvi, W., & Hillyard, S. A. (2007a). Spatial attention facilitates selection of illusory objects: evidence from event-related brain potentials. *Brain Research, 1139*, 143-152.

Martínez, A., Teder-Sälejärvi, W., Vazquez, M., Molholm, S., Foxe, J. J., Javitt, D. C., Di Russo, F., Worden, M. S., & Hillyard, S. A. (2006). Objects are highlighted by spatial attention. *Journal of Cognitive Neuroscience, 18*(2), 298-310.

Martinez-Trujillo, J. C, & Treue, S. (2004). Feature-based attention increases the selectivity of population responses in primate visual cortex. *Current Biology, 14*, 744-751.

Mathewson, K. E., Gratton, G., Fabiani, M., Beck, D. M., & Ro, T. (2009). To see or not to see: prestimulus alpha phase predicts visual awareness. *Journal of Neuroscience, 29*(9), 2725-2732.

McMains, S. A., Fehd, H. M;, Emmanouil, T.-A., & Kastner, S. (2007). Mechanisms of feature- and space-based attention: response modulation and baseline increases. *Journal of Neurophysiology, 98*, 2110-2121.

Molholm, S., Martinez, A., Shpaner, M., & Foxe, J. J. (2007). Object-based attention is multisensory: co-activation of an object's representations in ignored sensory modalities. *European Journal of neuroscience, 26*(2), 499-509.

Mulholland, T. B. (1965). Occurrence of the electroencephalographic alpha rhythm with eyes open. *Nature, 206*, 746.

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

Morgan, M. J., Ward, R. M., Castet, E. (1998). Visual search for a tilted target: tests of spatial uncertainty models. *Quarterly Journal of Experimental Psychology A: Human Experimental Psychology, 51*(2), 347-70.

Most, S. B., & Astur, R. S. (2007). Feature-based attentional set as a cause of traffic accidents. *Visual Cognition, 75*(2), 125-132.

Motter, B. C. (1993). Focal attention produces spatially selective processing in visual cortical areas V1, V2, and V4 in the presence of competing stimuli. *Journal of Neurophysiology, 70*(3), 909-919.

Murray, M. M., Foxe, D. M., Javitt, D. C., & Foxe, J. J. (2004). Setting boundaries: brain dynamics of modal and amodal illusory shape completion in humans. *Journal of Neuroscience, 24*(31), 6898-6903.

Murray, M. M., Imber, M. L., Javitt, D. C., & Foxe, J. J. (2006). Boundary completion is automatic and dissociable from shape discrimination. *Journal of Neuroscience, 26*(46), 12043-12054.

Murray, M. M., Wylie, G. R., Higgins, B. A., Javitt, D. C., Schroeder, C. E., & Foxe, J. J. (2002). The spatiotemporal dynamics of illusory contour processing: combined high-density electrical mapping, source analysis, and functional magnetic resonance imaging. *Journal of Neuroscience, 22*(12), 5055-5073.

Neely, J. H. (1977). Semantic priming and retrieval from lexical memory: roles of inhibitionless spreading activation and limited-capacity attention. *Journal of Experimental Psychology, 106*(3), 226-254.

O'Connell, R. G., Dockree, P. M., Robertson, I. H., Bellgrove, M. A., Foxe, J. J., & Kelly, S. P. (2009). Uncovering the neural signature of lapsing attention: electrophysiological signals predict errors up to 20s before they occur. *Journal of Neuroscience, 29*(26), 8604-8611.

O'Craven, K.M., Downing, P. E., & Kanwisher, N. (1999). fMRI evidence for objects as the units of attentional selection. *Nature*, *401*, 584-587.

Onton, J., & Makeig, S. (2006). Information-based modeling of event-related brain dynamics. *Progress in Brain Research*, *159,* 99-120.

Pavlovskaya, M., Ring, H., Groswasser, Z., & Hochstein, S. (2002). Searching with unilateral neglect. *Journal of Cognitive Neuroscience*, *14,* 745-756.

Peirce, C. S. (no date). Reflections on real and unreal objects. *Manuscript #966*. Quoted from http://en.wikiquote.org/wiki/Charles_Sanders_Peirce. Accessed 31 May 2011.

Perrott, D. R., Cisneros, J., McKinley, R. L., & D'Angelo, W. R. (1996). Aurally aided visual search under virtual and free-field listening conditions. *Human Factors, 38*(4), 702-715.

Perrott, D. R., Saberi, K., Brown, K., & Strybel, T. Z. (1990). Auditory psychomotor coordination and visual search performance. *Perception & Psychophysics, 48*(3), 214-226.

Petersen, S. E., Robinson, D. L., & Morris, J. D. (1987). Contributions of the pulvinar to visual spatial attention. *Neuropsychologia, 25*(1A), 97-105.

Pfurtscheller, G. (1992). Event-related synchronization (ERS): an electrophysiological correlate of cortical areas at rest. *Electroencephalography and Clinical Neurophysiology, 83*(1), 62-69.

Pfurtscheller, G., & Aranibar, A. (1977). Event-related cortical desynchronization detected by power measurements of scalp EEG. *Electroencephalography and Clinical Neurophysiology, 42*(6), 817-826.

Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology, 32*(1), 3-25.

Posner, M. I., & Cohen, Y. (1984). Components of attention. In H. Bouma & D. G. Bouwhuis (Eds.), *Attention and Performance X* (pp. 55-66). Hillsdale, NJ: Erlbaum.

Posner, M. I., & Petersen, S. E. (1990). The attention system of the human brain. *Annual Review of Neuroscience, 13*, 25-42.

Posner, M. I., Snyder, C. R., & Davidson, B. J. (1980). Attention and the detection of signals. *Journal of Experimental Psychology, 109*(2), 160-174.

Posner, M. I., Walker, J. A., Friedrich, F. J., & Rafal, R. D. (1984). Effects of parietal injury on covert orienting of attention. *Journal of Neuroscience*, *4*(7), 1863-1874.

Rahn, E., & Başar, E. (1993a). Prestimulus EEG-activity strongly influences the auditory evoked vertex response: a new method for selective averaging. *International Journal of Neuroscience*, *69*(1-4), 207-220.

Rahn, R., & Başar, E. (1993b). Enhancement of visual evoked potentials by stimulation during low prestimulus EEG stages. *International Journal of Neuroscience*, *72*(1-2), 123-136.

Rees, G., Frith, C. D., & Lavie, N. (1997). Modulating irrelevant motion perception by varying attentional load in an unrelated task. *Science*, *278*(5343), 1616-1619.

Reynolds, J. H., Chelazzi, L., & Desimone, R. (1999). Competitive mechanisms subserve attention in macaque areas V2 and V4. *Journal of Neuroscience*, *19*(5), 1736-1753.

Rihs, T. A., Michel, C. M., & Thut, G. (2007). Mechanisms of selective inhibition in visual spatial attention are indexed by alpha-band EEG synchronization. *European Journal of Neuroscience*, *25*(2), 603-610.

Rihs, T. A., Michel, C. M., & Thut, G. (2009). A bias for posterior alpha-band power suppression versus enhancement during shifting versus maintenance of attention. *NeuroImage*, *44*, 190-199.

Robinson, D. L. (1993). Functional contributions of the primate pulvinar. *Progress in Brain Research*, *95*, 371-380.

Romei, V., Rihs, T., Brodbeck, V., & Thut, G. (2008a). Resting electroencephalogram alpha-power over posterior sites indexes baseline visual cortex excitability. *Neuroreport, 19*(2), 203-208.

Romei, V., Brodbeck, V., Michel, C., Amedi, A., Pascual-Leone, A., & Thut, G. (2008b). Spontaneous fluctuations in posterior alpha-band EEG activity reflect variability in excitability of human visual areas. *Cerebral Cortex, 18*(9), 2010-2018.

Romei, V., Gross, J., & Thut, G. (2010). On the role of prestimulus alpha rhythms over occipito-parietal areas in visual input regulation: correlation or causation? *Journal of Neuroscience, 30*(25), 8692-8697.

Sauseng, P., Klimesch, W., Stadler, W., Schabus, M., Doppelmayr, M., Hanslmayr, S., Gruber, W. R., & Birbaumer, N. (2005). A shift of visual spatial attention is selectively associated with human EEG alpha activity. *European Journal of Neuroscience, 22*, 2917-2926.

Scherg, M., & Picton, T. W. (1991). Separation and identification of event-related potential components by brain electric source analysis. *Electroencephalography and Clinical Neurophysiology Supplement, 42,* 24-37.

Scherg, M., & Von Cramon, D. (1985). Two bilateral sources of the late AEP as identified by a spatio-temporal dipole model. *Electroencephalography and Clinical Neurophysiology, 62,* 32–44.

Schneider, W., & Fisk, A. D. (1982). Concurrent automatic and controlled visual search: can processing occur without resource cost? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 8*(4), 261-278.

Schneider, W., & Shiffrin, R. M. (1977). Controlled and automatic human information processing: I. Detection, search, and attention. *Psychological Review, 84*(1), 1-66.

Schoenfeld, M. A., Tempelmann, C, Martinez, A., Hopf, J. M., Settler, C, Heinze H. J., & Hillyard S. A. (2003). Dynamics of feature binding during object-selective attention. *Proceedings of the National Academy of Sciences USA, 100*, 11806-11811.

Schoenfeld, M. A., Hopf, J. M., Martinez, A., Mai, H. M., Sattler, C, Gasde, A., Heinze, H.-J., & Hillyard, S. A. (2007). Spatio-temporal analysis of feature-based attention. *Cerebral Cortex, 17*(10), 2468-2477.

Sehatpour, P., Molholm, S., Schwartz, T. H., Mahoney, J. R., Mehta, A. D., Javitt, D. C., Stanton, P. K., & Foxe, J. J. (2008). A human intracranial study of long-range oscillatory coherence across a frontal-occipital-hippocampal brain network during visual object processing. *Proceedings of the National Academy of Sciences USA, 105*(11), 4399-4404.

Shibata, T., Shimoyama, I., Ito, T., Abla, D., Iwasa, H., Koseki, K., Yamanouchi, N., Sato, T., Nakajima, Y. (1999). Attention changes the peak latency of the visual gamma-band oscillation of the EEG. *Neuroreport, 10*(6), 1167-1170.

Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending and a general theory. *Psychological Review, 84*(2), 127-190.

Shpaner, M., Murray, M. M., & Foxe, J. J. (2009). Early processing in the human lateral occipital complex is highly responsive to illusory contours but not to salient regions. *European Journal of Neuroscience, 30*(10), 2018-2028.

Siegel, M. H. (1965). Color discrimination as a function of exposure time. *Journal of the Optical Society of America, 55*(5), 566-568.

Simpson, G. V., Pflieger, M. E., Foxe, J. J., Ahlfors, S. P., Vaughan, H. G. Jr., Hrabe, J., Ilmoniemi, R. J., & Lantos, G. (1995). Dynamic neuroimaging of brain function. *Journal of Clinical Neurophysiology, 12*, 432–449.

Singer, W. (1993). Synchronization of cortical activity and its putative role in information processing and learning. *Annual Review of Physiology, 55*, 349-74.

Singer, W. (1999). Neuronal synchrony: a versatile code for the definition of relations? *Neuron, 24*(1), 49-65, 111-125.

Snyder, A. C., & Foxe, J. J. (2010). Anticipatory attentional suppression of visual features indexed by oscillatory alpha-band power increases: a high-density electrical mapping study. *The Journal of Neuroscience, 30*(11), 4024-4032.

Snyder, A. C., & Foxe, J. J. (2011). The countervailing forces of binding and selection in vision. *Cortex,*doi:10.1016/j.cortex.2011.05.003 [Epub.].

Sokolov, A., Lutzenberger, W., Marina, P., Hubert, P., Braun, C., & Birbaumer, N. (1999). Gamma-band MEG activity to coherent motion depends on task-driven attention. *Neuroreport, 10*(10), 1997-2000.

Spatz, W. B. (1979). The retino-geniculo-cortical pathway in Callithrix. II. The geniculo-cortical projection. *Experimental Brain Research, 36*(3), 401-410.

Spence, C., & Driver, J. (1996). Audiovisual links in endogenous covert spatial attention. *Journal of Experimental Psychology: Human Perception and Performance, 22*(4), 1005-1030.

Spence, C., & Driver, J. (1997a). Audiovisual links in exogenous covert spatial orienting. *Perception & Psychophysics, 59*(1), 1-22.

Spence, C., & Driver, J. (1997b). On measuring selective attention to an expected sensory modality. *Perception & Psychophysics, 59*(3), 389-403.

Steriade, M., & Deschênes, M. (1984). The thalamus as a neuronal oscillator. *Brain Research, 320*(1), 1-63.

Steriade, M., Deschênes, M., Domich, L., & Mulle, C. (1985). Abolition of spindle oscillations in thalamic neurons disconnected from nucleus reticularis thalami. *Journal of Neurophysiology, 54*(6), 1473-1497.

Steriade, M., Gloor, P., Llinás, R. R., Lopes de Silva, F. H., & Mesulam, M. M. (1990). Report of IFCN committee on basic mechanisms. Basic mechanisms of cerebral rhythmic activities. *Electroencephalography and Clinical Neurophysiology, 76*(6), 481-508.

Stoppel, C. M., Boehler, C. N., Sabelhaus, C, Heinze, H.-J., Hopf, J. M., & Schoenfeld, M. A. (2007). Neural mechanisms of spatial- and feature-based attention: a quantitative analysis. *Brain Research, 1181*, 51-60.

Tallon-Baudry, C., Bertrand, O., Delpuech, C., & Pernier, J. (1996). Stimulus specificity of phase-locked and non-phase-locked 40 Hz visual responses in human. *The Journal of Neuroscience, 16*(13), 4240-4249.

Tamler, E. (1970). Dissociation of color and form sense following stroke. *American Journal of Ophthalmology, 70*(1), 101-103.

Tanaka, K. (1996). Inferotemporal cortex and object vision. *Annual Review of Neuroscience, 19*, 109-139.

Thut, G., Nietzel, A., Brandt, S. A., & Pascual-Leone, A. (2006). Differential distribution of alpha-band electroencephalographic (EEG) activity over occipital cortex indexes visuo-spatial attention bias and predicts visual target detection. *Journal of Neuroscience*, *26*(37), 9494-9502.

Treisman, A. M. (2004). Psychological Issues in Selective Attention. In M. S. Gazzaniga (Ed.), *The Cognitive Neurosciences III* (pp. 529-44). Cambridge: MIT Press.

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

Treisman, A., & Schmidt, H. (1982). Illusory conjunctions in the perception of objects. *Cognitive Psychology, 14*(1), 107-141.

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

Vaina, L. M. (1994). Functional segregation of color and motion processing in the human visual cortex: clinical evidence. *Cerebral Cortex, 4*(5), 555-572.

Vallar, G., & Perani, D. (1986). The anatomy of unilateral neglect after right-hemisphere stroke lesions: a clinical/CT-scan correlation study of man. *Neuropsychologia, 24*(5), 609-622.

Van der Heijden, A. H. C. (1991). *Selective attention in vision.* London: Routledge.

van Dijk, H., Schoffelen, J. M., Oostenveld, R., & Jensen, O. (2008). Prestimulus oscillatory activity in the alpha band predicts visual discrimination ability. *Journal of Neuroscience, 28*(8), 1816-1823.

van Ede, F., de Lange, F., Jensen, O., & Maris, E. (2011). Orienting attention to an upcoming tactile event involves a spatially and temporally specific modulation of sensorimotor alpha- and beta-band oscillations. *Journal of Neuroscience, 31*(6), 2016-2024.

Vanduffel, W., Tootell, R. B., & Orban, G. A. (2000). Attention-dependent suppression of metabolic activity in the early stages of the macaque visual system. *Cerebral Cortex, 10*(2), 109-126.

Vanni, S., Revonsuo, A., & Hari, R. (1997). Modulation of the parieto-occipital alpha rhythm during object detection. *The Journal of Neuroscience, 17*(18), 7141-7147.

Versace, F., Robinson, J. D., Lam, C. Y., Minnix, J. A., Brown, V. L., Carter, B. L., Wetter, D. W., & Cinciripini, P. M. (2010). Cigarette cues capture smokers' attention: evidence from event-related potentials. *Psychophysiology, 47*(3), 435-411.

Wertheimer, M. (1923). Untersuchungen zur Lehre von der Gestalt II. *Psycologische Forschung, 4*, 301-350.

Wetherill, G. B., & Levitt, H. (1965). Sequential estimation of points on a psychometric function. *British Journal of Mathematical and Statistical Psychology, 18*, 1-10.

Whittingstall, K., Bartels, A., Singh, V., Kwon, S., & Logothetis, N. K. (2010). Integration of EEG source imaging and fMRI during continuous viewing of natural movies. *Magn. Reson. Imaging, 28*(8), 1135-1142.

Williamson, S. J., Kaufman, L., Lu, Z. L., Wang, J. Z., & Karron, D. (1997). Study of human occipital alpha rhythm: the alphon hypothesis and alpha suppression. *International Journal of Psychophysiology, 26*(1-3), 63-76.

Worden, M. S., Foxe, J. J., Wang, N., & Simpson, G. V. (2000). Anticipatory biasing of visuospatial attention indexed by retinotopically specific alpha-band electroencephalography increases over occipital cortex. *Journal of Neuroscience, 20*(6), RC63.

Wörgötter, F., Suder, K., Zhao, Y., Kerscher, N., Eysel, U. T., & Funke, K. (1998). State-dependent receptive-field restructuring in the visual cortex. *Nature, 396*(6707), 165-168.

Wyart, V., & Tallon-Baudry, C. (2009). How ongoing fluctuations in human visual cortex predict perceptual awareness: baseline shift versus decision bias. *Journal of Neuroscience, 29*(27), 8715-8725.

Wylie, G. R., Javitt, D. C., & Foxe, J. J. (2004). Don't think of a white bear: an fMRI investigation of the effects of sequential instructional sets on cortical activity in a task-switching paradigm. *Human Brain Mapping, 21*(4), 279-297.

Wylie, G. R., Javitt, D. C., & Foxe, J. J. (2006). Jumping the gun: is effective preparation contingent upon anticipatory activation in task-relevant neural circuitry? *Cerebral Cortex, 16*, 394-404.

Yamagishi, N., Goda, N., Callan, D. E., Anderson, S. J., & Kawato, M. (2005). Attentional shifts towards an expected visual target alter the level of alpha-band oscillatory activity in the human calcarine cortex. *Brain Research: Cognitive Brain Research, 25*(3), 799-809.

Yantis, S., & Jonides, J. (1984). Abrupt visual onsets and selective attention: evidence from visual search. *Journal of Experimental Psychology: Human Perception and Performance, 10,* 601-621.

Zanto, T. P., & Gazzaley, A. (2009). Neural suppression of irrelevant information underlies optimal working memory performance. *Journal of Neuroscience, 29*(10), 3059-3066.