Behavioral/Systems/Cognitive

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

Adam C. Snyder[1] and John J. Foxe[1,2]

[1]Program in Cognitive Neuroscience, Departments of Psychology and Biology, City College of the City University of New York, New York, New York 10031, and [2]The Cognitive Neurophysiology Laboratory, Nathan S. Kline Institute for Psychiatric Research, Program in Cognitive Neuroscience and Schizophrenia, Orangeburg, New York 10962

Retinotopically specific increases in alpha-band (~10 Hz) oscillatory power have been strongly implicated in the suppression of processing for irrelevant parts of the visual field during the deployment of visuospatial attention. Here, we asked whether this alpha suppression mechanism also plays a role in the nonspatial anticipatory biasing of feature-based attention. Visual word cues informed subjects what the task-relevant feature of an upcoming visual stimulus (S2) was, while high-density electroencephalographic recordings were acquired. We examined anticipatory oscillatory activity in the Cue-to-S2 interval (~2 s). Subjects were cued on a trial-by-trial basis to attend to either the color or direction of motion of an upcoming dot field array, and to respond when they detected that a subset of the dots differed from the majority along the target feature dimension. We used the features of color and motion, expressly because they have well known, spatially separated cortical processing areas, to distinguish shifts in alpha power over areas processing each feature. Alpha power from dorsal regions increased when motion was the irrelevant feature (i.e., color was cued), and alpha power from ventral regions increased when color was irrelevant. Thus, alpha-suppression mechanisms appear to operate during feature-based selection in much the same manner as has been shown for space-based attention.

Introduction

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 (Hillyard et al., 1998; McMains et al., 2007). Conversely, if a region of space is expected to be a locus of distracting events, processing of stimuli occurring there is attenuated (Rees et al., 1997; Hillyard et al., 1998). 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 (Luck et al., 1997; Foxe et al., 2005; 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 (~8–15 Hz) oscillatory power in the EEG (Worden et al., 2000; Kelly et al., 2005, 2006; Sauseng et al., 2005; Yamagishi et al., 2005; Thut et al., 2006; Rihs et al., 2007).

Based on the cellular physiology of similar oscillations in animals, it has been proposed that alpha might serve as a functional gating mechanism (Lopes da Silva, 1991; Foxe et al., 1998). In a direct and compelling test of this gating role for alpha activity, Romei et al. (2008) stimulated visual cortex with transcranial magnetic stimulation (TMS) while monitoring alpha power. They found that the probability 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 (Lopes da Silva, 1991; Lindgren et al., 1999), is implicated in several influential theories of attention (Posner and Petersen, 1990; LaBerge, 1997). 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 frontoparietal attention network, Capotosto et al. (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 charac-

terized with respect to spatial and intersensory attention (Foxe et al., 1998; Fu et al., 2001; Gomez-Ramirez et al., 2007). Attention can also be deployed to nonspatial visual features, such as color or motion parameters, facilitating the processing of subsequent stimuli incorporating the attended feature, independently of spatial location (Corbetta et al., 1991; Martinez-Trujillo and Treue, 2004; Most and Astur, 2007; Egner et al., 2008). Here, we asked whether the role of alpha-band oscillations is specific to spatial and intersensory attentional selection, or whether 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 spatio-temporal 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).

Materials and Methods

Participants. Twelve adults (9 male, two left-handed) aged 21–50 years (mean: 30.5 ± 8.2 years) 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 below, Stimuli). 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 before 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 used a variant of the common S1–S2 cuing paradigm (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 whether 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. Noninformative (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 (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
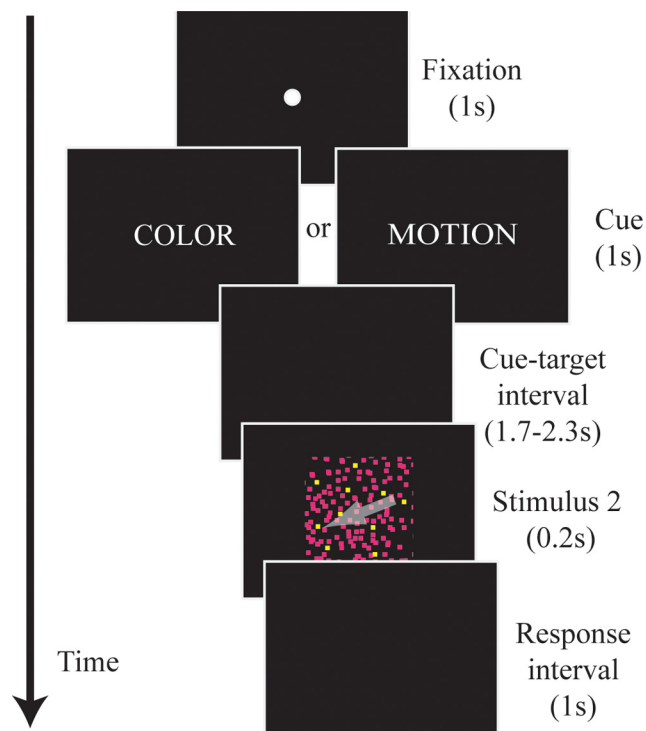

Figure 1. Schematic of task procedures. For each trial, subjects first viewed a fixation dot for 1 s, followed by a cue word in block capitals for 1 s. The cue (S1) was followed by an interval of 1.7–2.3 s with no stimulation. After the cue-target interval, the random dot stimulus (S2) was shown for 0.2 s, followed by a 1 s 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 is not to scale.

investigating a measure of attentional suppression (i.e., alpha). On the other hand, instructional cues [as in the study by 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 used a probabilistic cue to determine whether attention can be selectively used in our feature-based design as indicated by the standard reaction time measure, which was in fact the case. 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 75 Hz refresh rate. Trials began with a warning cue consisting of a white fixation dot on a black background for 1 s, followed by a cue word in white block capitals ("COLOR," "HUE," "MOTION," or "DIRECTION") for 1 s. 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 1.7–2.3 s (random and evenly distributed) during which only a black screen was displayed, the S2 was presented for 0.2 s (Fig. 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 18–36 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 et al., 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 above, Participants). 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 whether 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 nontargets 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 1 s response interval. Each subsequent trial began immediately following the response interval.

Before 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 and Levitt, 1965). After titration, subjects completed ten 10 min blocks (with self-paced breaks after every 12th trial).

EEG recording. Continuous EEG was acquired through the ActiveTwo BioSemi electrode system from 168 scalp electrodes, digitized at 512 Hz. 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 <1 Ω. 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 http://www.biosemi.com/faq/cms&drl.htm. During online data collection, signals were bandpass filtered between 0.1–100 Hz. Data were re-referenced offline to the average activity and downsampled to 32 Hz (see below, Independent component analysis). EEG data were processed using the FieldTrip toolbox (Donders Institute for Brain, Cognition and Behavior, Radboud University Nijmegen, Nijmegen, the Netherlands. See http://www.ru.nl/neuroimaging/fieldtrip) for MATLAB (The MathWorks).

Independent component analysis. 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 bandpass filter the scalp-recorded data at the frequency range of interest, and then estimate the inverse source solution of the filtered data (see supplement 1, available at www.jneurosci.org as supplemental material). 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 independent component analysis (ICA) can robustly separate artifacts from brain-related activity (Delorme et al., 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 are currently 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 (3×10^6 time-points $\times$ 168 channels), we first downsampled the timeseries to 32 Hz sampling rate for computational tractability. This preserves frequency information <16 Hz and thus does not impact the planned analysis in the 8–15 Hz alpha frequency band.

Temporal spectral evolution analysis. To examine the spatiotemporal dynamics of alpha-band amplitude in the cue-target interval, temporal spectral evolution (TSE) waveforms were derived by the following method. First, epochs time-locked to the cue (200 ms pre- to 3500 ms postcue) were bandpass filtered (third-order digital Butterworth zerophase, 8–15 Hz, 24 dB 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 (1.7–2.7 s postcue 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 of three SDs 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 (1.7–2.7 s postcue 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 (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 (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 6 and 7 mm, respectively, 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 used a non-parametric bootstrapping procedure. This approach has the advantage that it makes no

Table 1. Summary of components

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

For each subject, the number of components for which alpha power increased (or decreased) by at least 3 SDs 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.
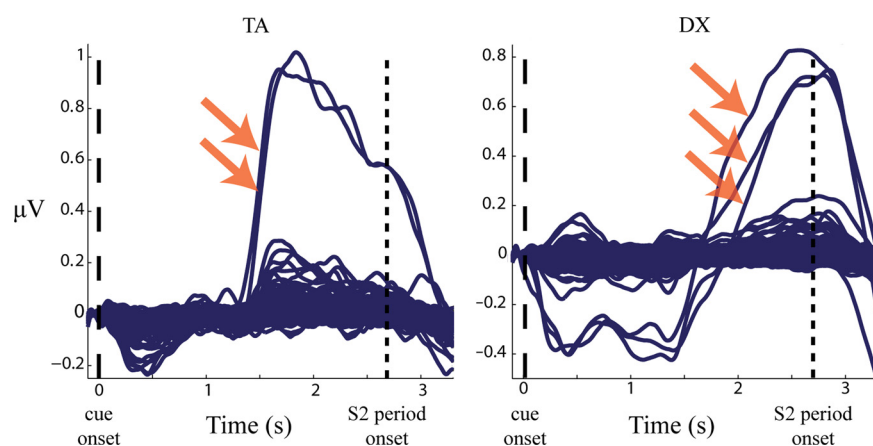

Figure 2. TSE waveforms for independent components. Sample data are shown for two subjects in the cue-target interval. Each trace is a TSE waveform showing the average time course of induced amplitude in the 8 –15 Hz alpha band for one independent component. Arrows indicate "alpha-reactive" components, which increase by more than three SDs in the final second of the preparatory period.

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 intergroup 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 (Kelly et al., 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 200 ms 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 0 ms 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 12 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 (supplement 2, available at www.jneurosci.org as supplemental material). 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 below, Behavior).

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 (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

with 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 values >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 (all r values $= 0.14$–0.24; all p values < 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 whether this could be related to their ability to make each type of discrimination. While performance was pretitrated 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 (Fig. 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 with 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.
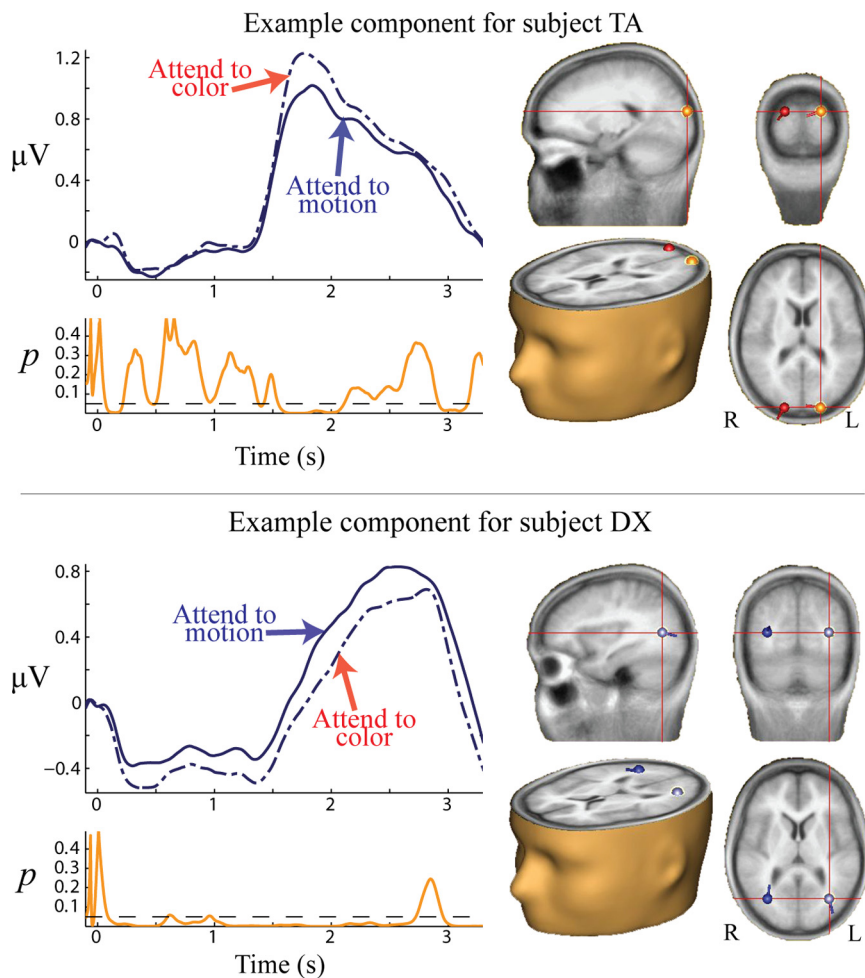

Figure 3. Feature-sensitive alpha-reactive components. Sample data are shown for two subjects in the cue-target interval. Left, TSE waveforms for attention to color (dashed) and motion (solid), with corresponding p values for the difference between the two conditions shown below in orange. For the p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

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 manner. Specifically, the decrease was greater when motion was cued. This
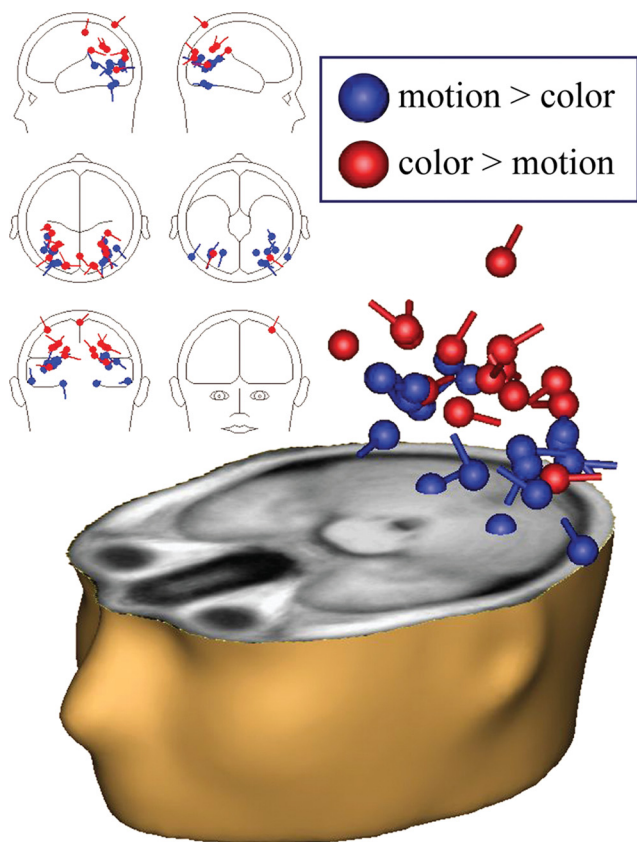
Figure 4. Source localizations for all feature-sensitive alpha-reactive components. Components with greater alpha power for attention to motion than for attention to color are plotted in blue and are generally localized to ventral visual stream regions. Components with greater alpha power for attention 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 subjects may have had more than one component that was feature sensitive.

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 (Hanslmayr et al., 2007; Romei et al., 2008). 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 toward 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 (2007) that could account for this latter difference in findings. First, of course, there are substantial differences in the nature of the tasks used. 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. Furthermore, 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

Table 2. Summary of component source localization

	Left Hemisphere				Right Hemisphere				
	x	y	z	VAF%	x	y	z	VAF%	Total 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

Positions are given in approximate Talairach coordinates. Only positively reactive components are listed. [†]VAF reported includes a third dipole.
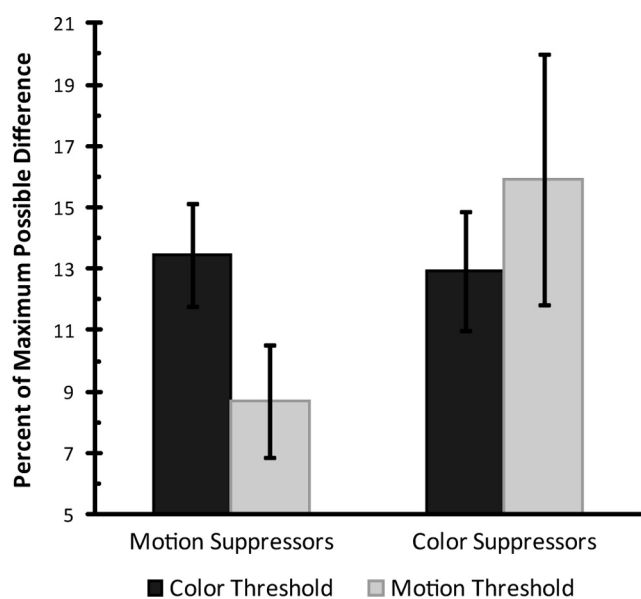

Figure 5. Motion direction and hue discrimination thresholds for motion and color suppressors. Units are the percentage of the maximum possible difference between feature values. Subjects having only components that increased alpha power when motion was attended relative to when color was attended (color 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.

(2007) 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 >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. How-

ever, many studies have shown that such relationships between alpha-band power and performance accuracy exist for spatial attention (Kelly et al., 2009; Wyart and Tallon-Baudry, 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., 2008; van Dijk et al., 2008; Mathewson et al., 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 et al. (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 et al., 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 compared with atten-

tion to the other feature, alpha-band power increases relative to baseline occurred regardless of which feature was cued. This observation could be related to the suppression observed by Martinez-Trujillo and Treue (2004) at the single-cell level.

A great deal of research has shown that there is a strong bias to attend to objects as wholes (Treisman, 2004; Blaser et al., 2000; Martínez 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 (Liu et al., 2003; McMains et al., 2007; Stoppel et 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.

References

Blaser E, Pylyshyn ZW, Holcombe AO (2000) Tracking an object through feature space. Nature 408:196–199.

Capotosto P, Babiloni C, Romani GL, Corbetta M (2009) Frontoparietal cortex controls spatial attention through modulation of anticipatory alpha rhythms. J Neurosci 29:5863–5872.

Corbetta M, Miezin FM, Dobmeyer S, Shulman GL, Petersen SE (1991) Selective and divided attention during visual discriminations of shape, color, and speed: functional anatomy by positron emission tomography. J Neurosci 11:2383–2402.

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 AM, Krauskopf J, Lennie P (1984) Chromatic mechanisms in lateral geniculate nucleus of macaque. J Physiol 357:241–265.

Egner T, Monti JM, Trittschuh EH, Wieneke CA, Hirsch J, Mesulam MM (2008) Neural integration of top-down spatial and feature-based information in visual search. J Neurosci 28:6141–6151.

Foxe JJ, Simpson GV, Ahlfors SP (1998) Parieto-occipital ~10 Hz activity reflects anticipatory state of visual attention mechanisms. Neuroreport 9:3929–3933.

Foxe JJ, Simpson GV, Ahlfors SP, Saron CD (2005) Biasing the brain's attentional set: I. Cue driven deployments of intersensory selective attention. Exp Brain Res 166:370–392.

Fu KM, Foxe JJ, Murray MM, Higgins BA, Javitt DC, Schroeder CE (2001) Attention-dependent suppression of distracter visual input can be cross-modally cued as indexed by anticipatory parieto-occipital alpha-band oscillations. Brain Res Cogn Brain Res 12:145–152.

Gomez-Ramirez M, Higgins BA, Rycroft JA, Owen GN, Mahoney J, Shpaner M, Foxe JJ (2007) The deployment of intersensory selective attention: a high-density electrical mapping study of the effects of theanine. Clin Neuropharmacol 30:25–38.

Graimann B, Pfurtscheller G (2006) Quantification and visualization of event-related changes in oscillatory brain activity in the time-frequency domain. Prog Brain Res 159:79–97.

Hanslmayr S, Aslan A, Staudigl T, Klimesch W, Herrmann CS, Bäuml KH (2007) Prestimulus oscillations predict visual perception performance between and within subjects. Neuroimage 37:1465–1473.

Hillyard SA, Vogel EK, Luck SJ (1998) Sensory gain control (amplification) as a mechanism of selective attention: electrophysiological and neuroimaging evidence. Philos Trans R Soc Lond B Biol Sci 353:1257–1270.

Hyvärinen A (1999) Fast and robust fixed-point algorithms for independent component analysis. IEEE Trans Neural Netw 10:626–634.

Jokisch D, Jensen O (2007) Modulation of gamma and alpha activity during a working memory task engaging the dorsal or ventral stream. J Neurosci 27:3244–3251.

Kelly SP, Lalor EC, Reilly RB, Foxe JJ (2005) Visual spatial attention tracking using high-density SSVEP data for independent brain-computer communication. IEEE Trans Neural Syst Rehabil Eng 13:172–178.

Kelly SP, Lalor EC, Reilly RB, Foxe JJ (2006) Increases in alpha oscillatory power reflect an active retinotopic mechanism for distracter suppression during sustained visuospatial attention. J Neurophysiol 95:3844–3851.

Kelly SP, Gomez-Ramirez M, Foxe JJ (2009) The strength of anticipatory spatial biasing predicts target discrimination at attended locations: a high-density electroencephalographic study. Eur J Neurosci 30:2224–2234.

LaBerge D (1997) Attention, awareness, and the triangular circuit. Conscious Cogn 6:149–181.

Lindgren KA, Larson CL, Schaefer SM, Abercrombie HC, Ward RT, Oakes TR, Holden JE, Perlman SB, Benca RM, Davidson RJ (1999) Thalamic metabolic rate predicts EEG alpha power in healthy control subjects but not in depressed patients. Biol Psychiatry 45:943–952.

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

Liu T, Slotnick SD, Serences JT, Yantis S (2003) Cortical mechanisms of feature-based attentional control. Cereb Cortex 13:1334–1343.

Lopes da Silva F (1991) Neural mechanisms underlying brain waves: from neural membranes to networks. Electroencephalogr Clin Neurophysiol 79:81–93.

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

Martínez A, Teder-Sälejärvi W, Vazquez M, Molholm S, Foxe JJ, Javitt DC, Di Russo F, Worden MS, Hillyard SA (2006) Objects are highlighted by spatial attention. J Cogn Neurosci 18:298–310.

Martinez-Trujillo JC, Treue S (2004) Feature-based attention increases the selectivity of population responses in primate visual cortex. Curr Biol 14:744–751.

Mathewson KE, Gratton G, Fabiani M, Beck DM, Ro T (2009) To see or not to see: prestimulus alpha phase predicts visual awareness. J Neurosci 29:2725–2732.

McMains SA, Fehd HM, Emmanouil TA, Kastner S (2007) Mechanisms of feature- and space-based attention: response modulation and baseline increases. J Neurophysiol 98:2110–2121.

Molholm S, Martinez A, Shpaner M, Foxe JJ (2007) Object-based attention is multisensory: co-activation of an object's representations in ignored sensory modalities. Eur J Neurosci 26:499–509.

Most SB, Astur RS (2007) Feature-based attentional set as a cause of traffic accidents. Vis Cogn 75:125–132.

O'Craven KM, Downing PE, 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. Prog Brain Res 159:99–120.

Posner MI (1980) Orienting of attention. Q J Exp Psychol 32:3–25.

Posner MI, Petersen SE (1990) The attention system of the human brain. Annu Rev Neurosci 13:25–42.

Rees G, Frith CD, Lavie N (1997) Modulating irrelevant motion perception by varying attentional load in an unrelated task. Science 278:1616–1619.

Rihs TA, Michel CM, Thut G (2007) Mechanisms of selective inhibition in visual spatial attention are indexed by alpha-band EEG synchronization. Eur J Neurosci 25:603–610.

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

Sauseng P, Klimesch W, Stadler W, Schabus M, Doppelmayr M, Hanslmayr S,

Gruber WR, Birbaumer N (2005) A shift of visual spatial attention is selectively associated with human EEG alpha activity. Eur J Neurosci 22:2917–2926.

Scherg M, Picton TW (1991) Separation and identification of event-related potential components by brain electric source analysis. Electroencephalogr Clin Neurophysiol [Suppl] 42:24–37.

Scherg M, Von Cramon D (1985) Two bilateral sources of the late AEP as identified by a spatio-temporal dipole model. Electroencephalogr Clin Neurophysiol 62:32–44.

Schoenfeld MA, Tempelmann C, Martinez A, Hopf JM, Sattler C, Heinze HJ, Hillyard SA (2003) Dynamics of feature binding during object-selective attention. Proc Natl Acad Sci U S A 100:11806–11811.

Schoenfeld MA, Hopf JM, Martinez A, Mai HM, Sattler C, Gasde A, Heinze HJ, Hillyard SA (2007) Spatio-temporal analysis of feature-based attention. Cereb Cortex 17:2468–2477.

Simpson GV, Pflieger ME, Foxe JJ, Ahlfors SP, Vaughan HG Jr, Hrabe J, Ilmoniemi RJ, Lantos G (1995) Dynamic neuroimaging of brain function. J Clin Neurophysiol 12:432–449.

Stoppel CM, Boehler CN, Sabelhaus C, Heinze HJ, Hopf JM, Schoenfeld MA (2007) Neural mechanisms of spatial- and feature-based attention: a quantitative analysis. Brain Res 1181:51–60.

Thut G, Nietzel A, Brandt SA, 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. J Neurosci 26:9494–9502.

Treisman AM (2004) Psychological issues in selective attention. In: The cognitive neurosciences III (Gazzaniga MS, ed), pp 529–544. Cambridge, MA: MIT.

Ungerleider LG, Mishkin M (1982) Two cortical visual systems. In: Analysis of visual behavior (Ingle DJ, Goodale MA, Mansfield RJW, eds), pp 549–586. Cambridge, MA: MIT.

van Dijk H, Schoffelen JM, Oostenveld R, Jensen O (2008) Prestimulus oscillatory activity in the alpha band predicts visual discrimination ability. J Neurosci 28:1816–1823.

Wetherill GB, Levitt H (1965) Sequential estimation of points on a psychometric function. Br J Math Stat Psychol 18:1–10.

Worden MS, Foxe JJ, Wang N, Simpson GV (2000) Anticipatory biasing of visuospatial attention indexed by retinotopically specific alpha-band electroencephalography increases over occipital cortex. J Neurosci 20:RC63.

Wyart V, Tallon-Baudry C (2009) How ongoing fluctuations in human visual cortex predict perceptual awareness: baseline shift versus decision bias. J Neurosci 29:8715–8725.

Wylie GR, Javitt DC, Foxe JJ (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. Hum Brain Mapp 21:279–297.

Yamagishi N, Goda N, Callan DE, Anderson SJ, Kawato M (2005) Attentional shifts towards an expected visual target alter the level of alpha-band oscillatory activity in the human calcarine cortex. Brain Res Cogn Brain Res 25:799–809.

Zanto TP, Gazzaley A (2009) Neural suppression of irrelevant information underlies optimal working memory performance. J Neurosci 29:3059–3066.