

Feature-Based Attentional Modulations in the Absence of Direct Visual Stimulation

John T. Serences^{1,2,4,*} and Geoffrey M. Boynton^{3,4}

¹Department of Cognitive Sciences

²Center for Cognitive Neuroscience, University of California, Irvine, Irvine, CA 92697-5100, USA

³Department of Psychology, University of Washington, Seattle, WA 98195-1525, USA

⁴The Salk Institute for Biological Studies, La Jolla, CA 92037, USA

*Correspondence: john.serences@uci.edu

DOI 10.1016/j.neuron.2007.06.015

SUMMARY

When faced with a crowded visual scene, observers must selectively attend to behaviorally relevant objects to avoid sensory overload. Often this selection process is guided by prior knowledge of a target-defining feature (e.g., the color red when looking for an apple), which enhances the firing rate of visual neurons that are selective for the attended feature. Here, we used functional magnetic resonance imaging and a pattern classification algorithm to predict the attentional state of human observers as they monitored a visual feature (one of two directions of motion). We find that feature-specific attention effects spread across the visual field—even to regions of the scene that do not contain a stimulus. This spread of feature-based attention to empty regions of space may facilitate the perception of behaviorally relevant stimuli by increasing sensitivity to attended features at all locations in the visual field.

INTRODUCTION

The human visual system has a limited processing capacity. Consequently, when multiple stimuli are simultaneously present in a scene, they must compete for cortical representation and access to awareness (Desimone and Duncan, 1995; Serences and Yantis, 2006). To resolve this competition, incoming sensory input is selectively filtered based on current behavioral goals so that relevant stimuli are processed more efficiently than irrelevant stimuli. An observer might attend to a particular region of space, providing a competitive advantage to stimuli presented at the selected location (Gandhi et al., 1999; Kastner et al., 1998; Moran and Desimone, 1985). Alternately, attended features (e.g., a color or direction of motion) will enjoy a competitive advantage over stimuli that do not express the attended feature, independent of their spatial location (Martinez-Trujillo and Treue, 2004; Melcher et al.,

2005; Saenz et al., 2002; Treue and Maunsell, 1996; Treue and Martinez Trujillo, 1999). Feature-based selection is thought to be especially important because we often know the defining features of a target (e.g., the pencil is yellow) without knowing its exact location (e.g., the pencil is somewhere on the desk).

Investigators have only recently begun to examine the neural basis of feature-based attention (Haenny et al., 1988; Motter, 1994). Treue and coworkers (Martinez-Trujillo and Treue, 2004; Treue and Maunsell, 1996; Treue and Martinez Trujillo, 1999) demonstrated that feature-based attention amplifies the response of a neuron when attention is directed to the neuron's preferred feature and suppresses the response when attention is directed to the neuron's nonpreferred feature (see also Boynton, 2005). This "feature-similarity gain" mechanism operates on the firing rate of all neurons tuned to the attended feature, even when the neuron is being driven by a stimulus that is outside the current focus of spatial attention (Bichot et al., 2005; Martinez-Trujillo and Treue, 2004; Saenz et al., 2002; Treue and Martinez Trujillo, 1999). In addition, if the multiplicative gain factor also boosts the baseline firing rate of a neuron ("activity gain;" see, e.g., Williford and Maunsell, 2006), then the model makes the untested prediction that feature-based attention should modulate the firing rate of a neuron that is not directly driven by a stimulus in its spatial receptive field.

In the present study, functional magnetic resonance imaging (fMRI) and an image classification algorithm were used to demonstrate direction-selective attentional modulations in human occipital and parietal cortex (or intraparietal cortex [IPS]), replicating and extending recent reports showing similar attention effects in early regions of visual cortex (Kamitani and Tong, 2005, 2006). In addition, feature-based attentional modulations spread to stimuli presented outside the focus of spatial attention, corroborating single-cell recording studies (Bichot et al., 2005; Treue and Martinez Trujillo, 1999) and establishing the feature-specific nature of these spatially global modulations in human observers (Saenz et al., 2002). Finally, feature-based attention spreads to unstimulated regions of the visual scene, which may facilitate visual search by automatically priming behaviorally relevant features simultaneously across all locations in the visual field.

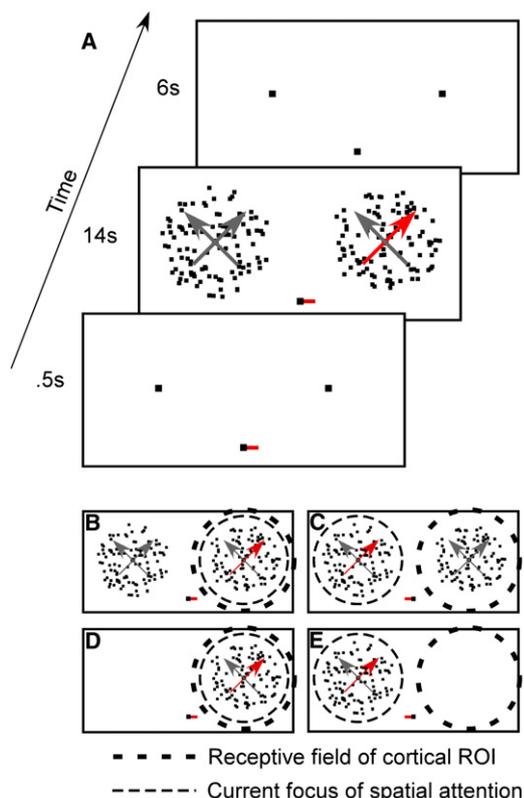


Figure 1. Behavioral Task

(A) Sequence of events on a trial where the observer was attending to 45° motion in the right stimulus aperture. In this sequence, stimuli were presented in both the attended and the unattended locations; however, on one-half of the trials, no stimulus was presented at the unattended location. One-half of the dots in each stimulus aperture moved at 45°, and the other half moved at 135°, for the duration of the 14 s presentation period (arrows shown in figure were not present in the actual display). The central attention cue indicated both the location of the to-be-attended target stimulus (by pointing left or right) and the to-be-attended direction of motion (e.g., red = attend 45°; green = attend 135°). Targets were defined as a brief slowing of the dots at the attended location that moved in the currently attended direction (45° in this figure); distractors were defined as a brief slowing of the dots at the attended location that moved in the unattended direction (135° in this figure).

(B–E) The different stimulus configurations with respect to a cortical region of interest (ROI, e.g., left hMT+) receiving input from the right visual field. (B) The focus of spatial attention is in the receptive field of the ROI under consideration, and there are two stimulus apertures present. (C) The focus of spatial attention is outside the receptive field of the ROI under consideration, and there are two stimulus apertures present. (D) The focus of spatial attention is in the receptive field of the ROI under consideration, but there is only one stimulus aperture present. (E) The focus of spatial attention is outside the receptive field of the ROI and there is only one stimulus present in the visual field. In this case, the ROI receives input from an unstimulated region of space.

RESULTS

The behavioral task is depicted in Figure 1. Observers maintained gaze on a central fixation point; at the onset

of each trial, a small attention cue indicated whether to attend to the left or right stimulus aperture, and it also indicated the direction of motion to be monitored within the attended aperture (i.e., either 45° or 135° motion). Observers pressed a button when the dots moving in the currently attended direction slowed (a target event), which occurred twice on each trial. No button press was required when the dots moving in the unattended direction slowed (a distractor event), which also occurred twice on each trial. On half of the trials, only the spatially cued aperture contained moving dots; on the remaining trials, dot fields were also presented on the unattended side of space. The speed of the moving dots presented on the unattended side of space (when present) remained constant throughout a trial. Each scan consisted of 24 trials interleaved with blank 6 s intertrial intervals. Observers were able to reliably discriminate target from distractor events (mean d' , \pm SEM with one stimulus aperture: 2.58 ± 0.33 , and with two stimulus apertures: 2.38 ± 0.30 , not significantly different).

Feature-Based Attentional Modulations in Visual Cortex

Visual features fall into superordinate categories such as motion or color, and can be further divided into subordinate categories such as specific directions of motion or specific colors. Traditionally, fMRI studies have been restricted to the superordinate level of analysis because the blood oxygenation level-dependent (BOLD) response measured with fMRI is spatially imprecise with respect to the topology of subordinate-level selectivity within visual cortex. For example, motion-selective region MT contains direction-selective columns of neurons, with a 180° array of columns spanning approximately 0.5 mm of cortex (Albright et al., 1984). Thus, a single fMRI voxel, measuring approximately 3 mm³ in the present study, should contain columns selective for many different directions. Averaging across all voxels in MT would nullify any directional selectivity in the BOLD response.

However, Kamitani and Tong (2005, 2006) recently circumvented this limitation by using multivariate pattern classification methods to measure subordinate-level feature-based representations (see also Cox and Savoy, 2003; Haxby et al., 2001; Haynes and Rees, 2005, 2006; Mitchell et al., 2003; Norman et al., 2006; Peelen and Downing, 2007). The method assumes that a preponderance of neurons preferring a particular feature might happen to be sampled within a single fMRI voxel, giving rise to a small but detectable feature-selective response bias. By considering the pattern of activity across many weakly selective voxels, Kamitani and Tong (2005, 2006) were able to predict the orientation or the direction of motion an observer was attending. Haynes and Rees (2005) used a similar procedure to predict the orientation of a stimulus rendered “invisible” via a combination of forward and backward masking (Macknik’s “standing wave of invisibility;” see Macknik and Livingstone, 1998).

Here, we used a pattern classification analysis to examine motion-specific attentional modulations in visually responsive regions of occipital, parietal, and frontal cortex (see [Experimental Procedures](#) and [Haynes and Rees, 2005](#)). All neuroimaging analyses were based on measurements of the BOLD response made within regions of interest (ROIs) that were independently identified using a functional localizer task that isolated voxels responding more strongly to stimuli in one visual hemifield compared with the other. We exploited the fact that ROIs in left visual cortex receive input primarily from stimuli in the right visual field, whereas ROIs in right visual cortex receive input primarily from stimuli in the left visual field. This contralateral stimulus-to-cortex mapping allowed us to measure BOLD responses evoked by attended and ignored stimulus apertures on the same trial because the two apertures projected to visual areas in opposite cortical hemispheres. The qualitative pattern of responses was similar across corresponding ROIs (e.g., left and right human MT+, or hMT+, which is thought to be the homolog to monkey regions MT and MST), so data were collapsed across hemispheres within each observer.

Each observer participated in either seven or eight scans of the main experimental task shown in [Figure 1](#). To classify the observer's attentional state (e.g., attending 45° versus 135° motion), we first computed an activation vector indexing the magnitude of the BOLD signal in each voxel within a given ROI at each time point from 6 s to 18 s poststimulus. We then grouped these activation vectors into three bins based on the following stimulus/attention configurations with respect to each ROI: trials on which the observer was attending to a contralateral stimulus ([Figures 1B](#) and [1D](#)), trials on which the observer was attending to an ipsilateral stimulus while there was an unattended stimulus in the contralateral field ([Figure 1C](#)), and trials on which the observer was attending to an ipsilateral stimulus while no stimulus was present in the contralateral visual field ([Figure 1E](#)). In all cases, the goal was to infer the currently attended direction of motion based on the activation vectors from each bin.

To carry out the classification in a given stimulus/attention configuration, activation vectors from all but one of the scans were averaged to form a “training” vector, and activation vectors from the remaining scan were averaged to form a “test” vector (a scan refers to an ~8 min data acquisition block, so training and test vectors were independent). To classify the currently attended direction of motion, each test activation vector was classified based on the mean training activation vector that it most closely resembled (see [Experimental Procedures](#)). The accuracy of classification was validated using a “hold-one-scan out” procedure, where data from each scan served as a test set in turn. Since the classification procedure was used to determine if the observers were attending to either 45° or 135° motion on any given trial, chance classification was always 50%.

[Figure 2A](#) shows classification accuracy within motion-selective regions V3A ([Braddick et al., 2001](#); [Orban et al.,](#)

[2003](#); [Tootell et al., 1997](#)) and hMT+ when the currently attended target stimulus was within the receptive field of the ROI under consideration (data collapsed across the stimulus configurations are depicted in [Figures 1B](#) and [1D](#); see [Figure S1](#) in the [Supplemental Data](#), available with this article online, for each condition separately). In V3A and hMT+, classification accuracy was above chance when all 50 voxels were considered, and the qualitative pattern of results was similar based on only the 30 most spatially selective voxels ([Figure S2](#)). As in other studies (e.g., the first figure in [Haynes and Rees, 2005](#)), classification accuracy was sometimes above chance when only a few voxels were considered because the Mahalanobis distance (md) is sensitive to differences in the pattern of activation across an ROI and differences in mean activation levels between conditions. Thus, differences in the mean activation level across the first several voxels supported above-chance classification of the attended direction (however, similar results were obtained when the mean of each activation vector was explicitly removed before pattern classification; see [Figure S3](#)).

This feature-based attentional modulation was not driven by sensory differences in the display because target events (a slowing of the dots moving in the currently attended direction) and distractor events (a slowing of the dots moving in the ignored direction) occurred with equal frequency on each trial. Spatial attention cannot account for the results because the moving stimuli were spatially interleaved, and a “spotlight” of attention would cover dots moving in both directions. Therefore, the successful classification of motion direction depicted in [Figure 2A](#) must reflect systematic changes in the pattern of activation across V3A and hMT+ induced by feature-based attention. [Figure 2B](#) depicts asymptotic classification accuracy—defined as classification accuracy with all 50 voxels considered—from all visual areas identified. Collapsed across visual areas, classification accuracy was significantly above chance [paired t test against chance, $t(9) = 5.84$, $p < 0.0005$]. Note that ROIs in IPS and frontal eye field (FEF) were only successfully identified in a subset of hemispheres across the ten observers (see [Experimental Procedures](#)).

If feature-based attention spreads automatically to other stimuli in the visual field ([Bichot et al., 2005](#); [Martinez-Trujillo and Treue, 2004](#); [Saenz et al., 2002](#); [Treue and Martinez Trujillo, 1999](#)), then the pattern of activation evoked by the unattended stimulus aperture should also be modulated by the current state of feature-based attention. [Figure 2C](#) shows classification accuracy based on responses in areas V3A and hMT+ when an ignored stimulus was in the contralateral visual field (as in [Figure 1C](#)). Even though the stimulus was completely irrelevant to the task, the activation patterns predicted the currently attended direction of motion, and similar results were obtained in other regions of occipital and parietal cortex [[Figure 2D](#), paired t test against chance collapsed across visual areas, $t(9) = 2.9$, $p < 0.025$]. Previous human neuroimaging studies have shown a general increase in the response to

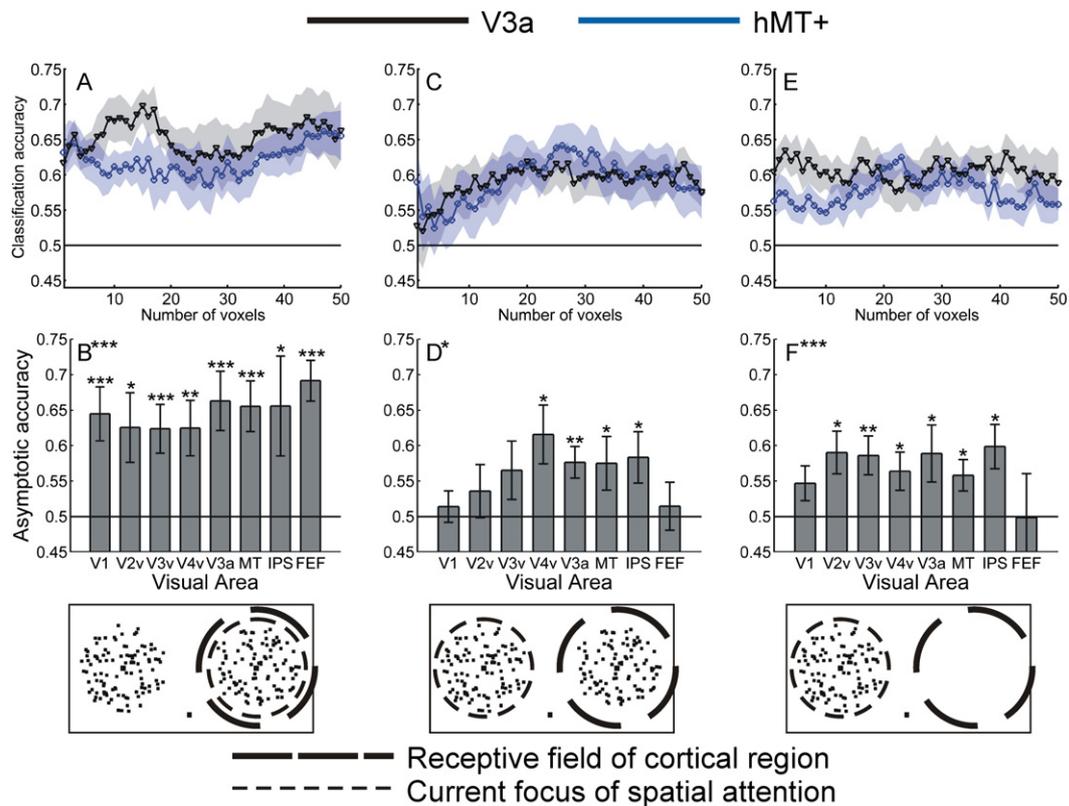


Figure 2. Pattern Classification Accuracy

(Top row) Accuracy of classifying the currently attended direction of motion as a function of the number of voxels included in the pattern analysis (shaded regions indicate ± 1 SEM across observers). (Bottom row) Asymptotic classification accuracy (defined as accuracy with all 50 voxels considered) in each of the visual ROIs (error bars, \pm SEM across observers). (A and B) Classification accuracy based on responses in ROIs contralateral to the focus of spatial attention (Figures 1B and 1D). (C and D) Classification accuracy based on responses to an ignored stimulus (Figure 1C). (E and F) Classification accuracy based on the responses to an unstimulated region of space (Figure 1E). * $p < 0.05$, ** $p < 0.01$, *** $p < 0.005$ based on planned repeated-measures t tests against chance (the asterisks by the panel labels in [B], [D], and [F] indicate significance of t tests computed on data collapsed across all visual areas).

ignored stimuli that share a feature with an attended stimulus (Saenz et al., 2002). The current data support the notion that these increased responses are due to response modulations within neurons selective for the attended feature.

If feature-based attention operates by enhancing baseline firing rates, then we reasoned that it might also modulate the pattern of activation within an ROI even in the absence of direct sensory stimulation (as in Figure 1E). Figure 2E shows classification accuracy based on responses in V3A and hMT+ when the contralateral visual field was unstimulated, and the pattern of activation within most ROIs predicted the currently attended direction of motion, even in the absence of direct sensory input [Figure 2F, paired t test against chance collapsed across visual areas, $t(9) = 5.29$, $p < 0.0005$]. These data provide evidence that feature-based attentional modulations spread to unstimulated regions of the visual field.

Since classification was carried out separately for each stimulus/attention configuration, the present results do

not establish that the pattern of activation observed in the absence of direct visual stimulation is the same as the pattern observed when an actual stimulus is present in the receptive field of an ROI. We explicitly tested this possibility by evaluating classification accuracy when the training set was based on attended contralateral motion and the test set was based on either an unattended motion stimulus (Figure 1C) or an unstimulated region of the visual field (Figure 1E). While classification accuracy for an unattended stimulus was slightly above chance in some visual areas (although nonsignificant overall), classification accuracy based on responses to an unstimulated region of the visual field was at chance. Thus, our results show that although feature-based attention induces a systematic modulation of the pattern of activation across an ROI even in the absence of direct stimulation (Figure 2), the activation patterns are not necessarily the same as the patterns observed when a stimulus is driving the response. We speculate that this lack of generalization may be due to large sensory differences between the

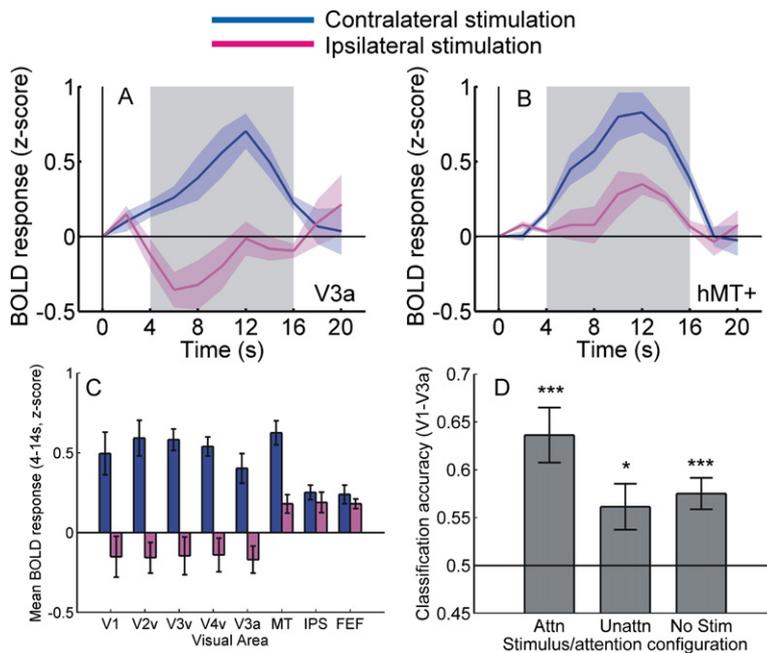


Figure 3. Spatial Selectivity of Voxels in Each ROI

The mean BOLD response during the spatial selectivity control experiment in which epochs of contralateral and ipsilateral peripheral stimulation were compared to epochs of no peripheral stimulation. All estimates are computed from the same 50 voxels that were selected for use in the main attention experiment.

(A and B) The time-windowed average time course of the BOLD response in V3A and hMT+ following the onset of contralateral and ipsilateral stimulation.

(C) Data from each visual area (bar plots show the mean response collapsed over 4 s to 16 s poststimulus, marked by the shaded region in [A] and [B]).

(D) Mean classification accuracy in each of the conditions depicted in Figures 2B, 2D, and 2F, collapsed across highly spatially selective regions V1, V2v, V3v, V4v, and V3A.

Error bars, \pm SEM across observers.

stimulus-present and stimulus-absent conditions. However, it is also possible that the pattern of attentional modulations in the absence of a stimulus may be qualitatively distinct from the pattern induced by an attended stimulus.

An alternate account of our results holds that above-chance classification accuracy might be observed in all conditions if the observer foveates the currently attended stimulus. However, eye position was monitored in several observers during scanning and the mean gaze position deviated less than 0.5° of visual angle during epochs of attention to the left and right sides of space (approximately the size of the central attention cue). Such small deviations in eye position are unlikely to significantly influence the data, given that the stimuli were located approximately $\pm 6^\circ$ in the periphery. Nor did we find any systematic differences in eye position related to the currently attended direction of motion (45° versus 135° , Figure S4).

Because we are arguing that feature-based attention spreads to unstimulated regions of the visual field, it is important to demonstrate that the spatial receptive fields of the selected voxels did not also encompass the currently attended ipsilateral stimulus. Four of the original observers were scanned in a control experiment in which 12 s epochs of left or right peripheral visual stimulation were interleaved with occasional 12 s epochs of passive fixation (see Experimental Procedures). As depicted in Figure 3 (see also Figure S5), the voxels within V1, V2v, V3v, V4v, and V3A that were used in the main attention experiment (i.e., Figure 2) were highly selective for epochs of contralateral stimulation compared with epochs of ipsilateral stimulation (where epochs of passive fixation formed the baseline). On the other hand, reliable positive responses in hMT+, IPS, and FEF were observed during

periods of ipsilateral stimulation compared with baseline (Ben Hamed et al., 2001; Tootell et al., 1998). However, our general conclusion that feature-based attention spreads to unstimulated regions of the visual field stands on the basis of data from early visual areas that exhibit a high degree of spatial selectivity (V1–V3A, Figure 3D). Of particular importance is the observation that area V3A, which is strongly motion selective in humans (Bradick et al., 2001; Orban et al., 2003; Tootell et al., 1997), shows both a spatially lateralized response and significant classification accuracy in the absence of direct visual stimulation.

We next tested the possibility that feature-based modulations in the absence of direct stimulation spread to regions of space beyond the defined stimulus apertures. We restricted our analysis to early regions V1, V2v, and V3v because (1) retinotopy is highly preserved (i.e., adjacent voxels respond to adjacent regions of the visual field), and (2) contamination by weak ipsilateral signals should be minimal (see, e.g., Tootell et al., 1998). First, we defined a “neighborhood” around each voxel in a given visual area (mean neighborhood size in voxels, \pm SEM: 45 ± 14 , see Experimental Procedures). To compute a “location selectivity index” within each neighborhood, we subtracted the mean response to ipsilateral stimuli from the mean response to contralateral stimuli during the functional localizer scans (Figure 4A). Since voxels in V1, V2v, and V3v should have receptive fields centered in the upper contralateral visual field, this index should be high for voxels that respond robustly to the contralateral stimulus aperture, and near zero for voxels that respond weakly to the contralateral aperture. The value of the selectivity index was assigned to the voxel at the center of the currently considered neighborhood, and this procedure was iterated until

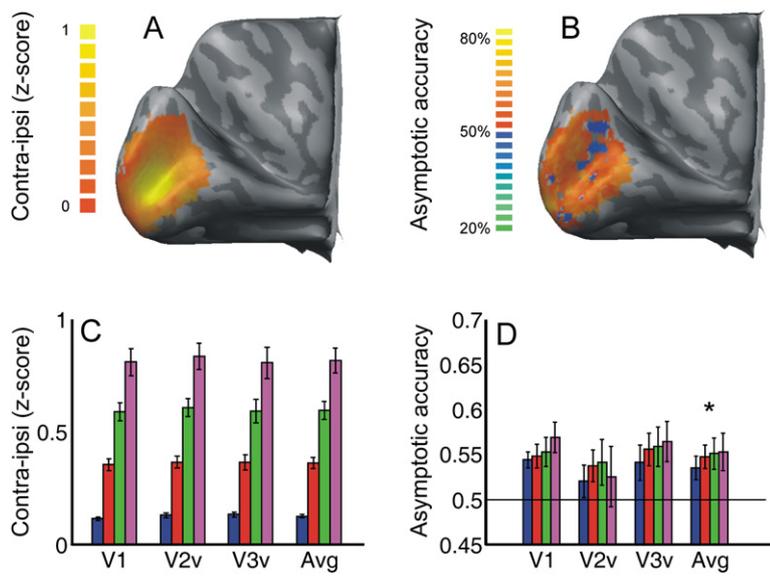


Figure 4. Spread of Attention beyond Stimulus Apertures

(A) Segmented and inflated left occipital lobe of a single observer depicting the stimulus selectivity index computed across V1 (see text and Experimental Procedures). Voxels rendered in yellow were highly selective for stimuli in the contralateral visual field, and voxels rendered in red were less selective. The bright yellow patch in the middle of V1 corresponds to the region from this observer that contained the most spatially selective voxels, which were used in the classification analyses shown in Figure 2.

(B) Classification accuracy for each voxel in the same observer in the absence of direct visual stimulation, based on the pattern of responses within that voxel's neighborhood (see text). Note that most regions of V1, even those not responding strongly to the contralateral stimulus aperture, correctly classified the attended direction >50% of the time. However, some subregions in this observer did not (shown in blue).

(C) Mean spatial selectivity index for voxels in each of four bins based on a quartile analysis. (D) Classification accuracy of voxels in each of the four selectivity bins shown in (C). (Accuracy averaged across V1, V2v, and V3v was significantly above chance in all four bins; * $p = 0.02$, * $p = 0.004$, * $p = 0.01$, * $p = 0.03$, respectively.) Error bars, \pm SEM across observers.

the spatial selectivity of each voxel was estimated. Next, we computed classification accuracy based on the pattern of responses across each neighborhood (Figure 4B). Finally, we sorted the voxels into four bins on the basis of their location selectivity (Figure 4C) and computed the average classification accuracy across all voxels falling into each of these bins (Figure 4D). Even voxels that were not highly responsive to the locations occupied by the stimulus apertures classified the currently attended feature. Similar results were obtained from the four observers that participated in the spatial selectivity control experiment where the location preference of each voxel was more directly estimated by comparing epochs of contralateral stimulation to passive fixation (Figure S6).

To determine if feature-based attention also spreads to regions of the lower visual field, where stimuli were never presented in the present experiment, we used the same neighborhood method to compute classification accuracy across visual areas V2d and V3d in each observer because neurons in these regions respond primarily to locations in the lower visual field. Again, we observed above-chance classification accuracy in each region (Figure 5), providing additional support for the hypothesis that feature-based modulations spread across the entire visual field.

To complement our ROI-based approach, a group random-effects analysis was carried out to evaluate classification accuracy across all of occipital cortex in the absence of direct visual stimulation. Using the neighborhood classification method described above, a sphere was defined around each voxel in the occipital cortex of an observer, and classification accuracy was computed based on the pattern of activation across the sphere.

The resulting single-observer classification accuracy maps were then standardized into Talairach space before averaging (Talairach and Tournoux, 1988). Significant classification accuracy was observed in areas near V3A and hMT+, as well as in several regions in striate and extrastriate visual cortex (Figure 6 and Table 1). While this group analysis lacks sensitivity because the visual areas of each observer are not perfectly aligned, the results generally confirm our ROI analyses and rule out the trivial possibility that classification accuracy is above chance for all voxels inside (or outside) the brain (see Experimental Procedures for additional details).

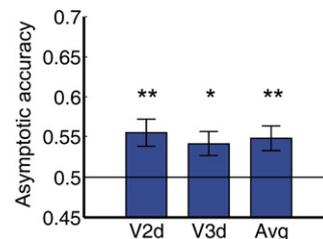


Figure 5. Spread of Attention to Lower Visual Field

Classification accuracy for dorsal occipital visual areas V2d and V3d, which respond primarily to locations in the lower visual field, in the absence of direct visual stimulation. Classification accuracy for each voxel was first computed for a neighborhood surrounding each voxel and then averaged across all voxels in each region. Data were not sorted based on stimulus selectivity (as in Figure 4) because voxels in these regions exhibit a generally poor response to stimuli in the upper visual field. * $p < 0.025$, ** $p < 0.01$. Error bars, \pm SEM across observers.

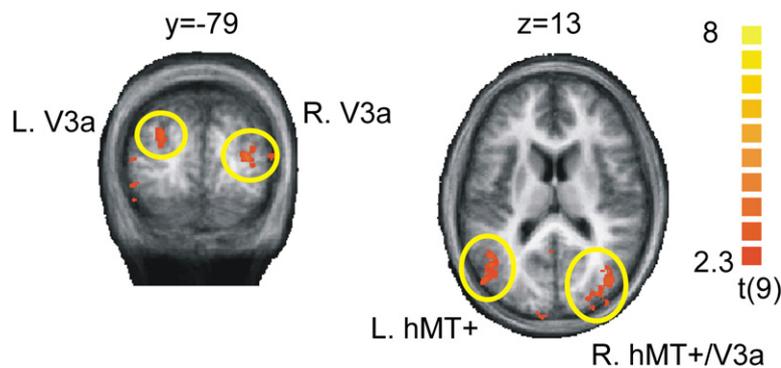


Figure 6. Group Random-Effects Analysis

One coronal and one transverse slice showing regions in the vicinity of V3A and hMT+ that exhibited above-chance classification accuracy in a group random-effects analysis. Brain images were generated by averaging anatomical scans across all subjects, and y and z coordinates are based on the atlas space of Talairach and Tournoux (1988) (images shown in neurological convention, left on left).

DISCUSSION

Here we show that feature-based attention spreads to stimuli presented outside the focus of spatial attention and to unstimulated regions of the visual field. The spatially global nature of feature-based attention may act to heighten sensitivity to relevant features across the visual field, a result predicted by models of visual search that posit an initial parallel stage of processing in which behaviorally relevant features are marked as having a high priority (e.g., Wolfe, 1994). In addition, the ballistic spread of feature-based attention across the visual field explains why behaviorally relevant features have a tendency to capture attention even when they are presented in an unexpected location (termed “contingent capture;” see Folk et al., 2002; Serences et al., 2005).

The spread of feature-based attention to stimuli outside the focus of spatial attention has been previously reported (Martinez-Trujillo and Treue, 2004; Saenz et al., 2002; Treue and Martinez Trujillo, 1999). However, in previous experiments the unattended aperture contained a single field of dots moving in either the attended or the unattended direction. While the investigators demonstrated that spatial attention was unlikely to play a significant role (Saenz et al., 2002, 2003), it remains possible that

observers may have covertly shifted spatial attention toward the to-be-ignored aperture more frequently when it contained an attended direction of motion than when it contained an unattended direction of motion. It is similarly possible that spatial attention was split between the two apertures more evenly when the to-be-ignored stimulus expressed the attended feature (Awh and Pashler, 2000; McMains and Somers, 2004). However, in our paradigm, when a stimulus was present in the ignored aperture, it always contained two overlapping dot fields moving in both the attended and ignored direction. Any spatial attention shift toward the ignored aperture would uniformly boost the response evoked by dots moving in both directions, instead of the observed direction-specific modulations reported in Figure 2. In addition, the observation of feature-specific modulations in the absence of direct visual stimulation cannot be attributed to spatial attention because the observer would have no reason to direct spatial attention toward a blank visual field. Together, these data strongly argue for the existence of a global feature-based attentional mechanism, and demonstrate that the observed modulations do not simply reflect spatial shifts of attention driven by feature similarity.

One concern is that neurons within a given ROI might have spatial receptive fields large enough to encompass

Table 1. Regions Exhibiting Significant Classification Accuracy

| General Extent of Region | Mean (x, y, z) | SEM (x, y, z) | Volume (ml) | t(9)* |
|------------------------------|----------------|---------------|-------------|-------|
| Left V1/dorsal extrastriate | (-7, -92, 7.4) | (5, 6, 5) | 1.02 | 2.9 |
| Left V1/ventral extrastriate | (-5, -57, 4) | (5, 8, 5) | 1.21 | 2.9 |
| Left V3A/posterior IPS | (-27, -69, 30) | (6, 7, 9) | 4.1 | 3.0 |
| Left hMT+ | (-45, -63, 15) | (4, 7, 4) | 3.2 | 3.1 |
| Right V1/dorsal extrastriate | (24, -87, 9) | (7, 6, 7) | 3.3 | 3.0 |
| Right ventral extrastriate | (8, -69, -12) | (6, 7, 6) | 1.1 | 2.8 |
| Right V3A/posterior IPS | (27, -64, 38) | (9, 9, 5) | 3.3 | 2.9 |
| Right hMT+ | (37, -73, 6) | (6, 7, 7) | 4.1 | 3.1 |

All regions exhibiting significant classification accuracy in the group random-effects analysis depicted in Figure 6. Mean coordinates (\pm SEM) are based on the atlas space of Talairach and Tournoux (1988), and t values reflect averages across all voxels in the cluster (* $p < 0.025$).

both stimulus apertures. If this were the case, a stimulus presented in the attended aperture would project to both contralateral and ipsilateral ROIs. While some neurons in higher-order visual areas such as V3A, hMT+, IPS, and FEF have receptive fields large enough to encompass both stimulus locations (Ben Hamed et al., 2001; Desimone and Ungerleider, 1986; Tootell et al., 1998), three arguments can be made against this general interpretation of the data. First, all pattern classification analyses were performed on the 50 most spatially selective voxels within each ROI as determined by independent functional localizer scans. Second, feature-specific attentional modulations were observed in regions such as V2v that are known to have small spatial receptive fields (e.g., $\sim 2^\circ$) centered primarily in the contralateral visual field (Desimone and Ungerleider, 1986). Finally, a control experiment comparing epochs of contralateral and ipsilateral stimulation with a low-level passive fixation baseline condition confirmed that the selected voxels within most ROIs (V1–V3A) had spatial receptive fields restricted to the contralateral hemifield. We do not dispute the existence of neurons (voxels) within many of these visual areas that respond to stimuli on both sides of fixation, especially in hMT+, IPS, and FEF, which were confirmed to respond to ipsilateral stimuli in our control experiment. However, we do argue that the selected voxels from most ROIs in the present study were responsive primarily to stimuli presented in the contralateral visual field.

The relatively high classification accuracy achieved based on responses to an unstimulated region of the visual field is surprising, given the relatively low baseline firing rate of the neurons in this condition (compare Figures 2B and 2F). Thus, feature-based attentional modulations would be occurring near the level of background noise, and one might expect these signals to be relatively weak. We hypothesize that the relatively high classification accuracy observed in the absence of direct visual stimulation might be due to the fact that visually responsive neurons are typically selective for a conjunction of several stimulus properties (e.g., spatial frequency and stimulus speed). Thus, when a stimulus is physically present within the receptive field of an ROI, neurons that are jointly tuned to the attended direction of motion *and* the other incidental features of the moving dots (size, speed, etc.) might dominate the response pattern. In contrast, when no stimulus is present within the receptive field, all neurons tuned to the currently attended direction of motion may undergo a systematic gain change. Since the classification algorithm takes into account the distributed pattern of activity across multiple voxels within each ROI, increasing the number of systematically modulated neurons might in turn support higher classification accuracy. Alternatively, the relatively high classification accuracy in the absence of direct stimulation may be the result of some nonlinearity in the BOLD response that makes feature-based attentional modulations easier to detect when baseline activity is lower.

A recent behavioral study provides indirect evidence supporting the present observation that feature-based attention spreads into empty regions of space (Arman et al., 2006). The study showed that attending to a single motion stimulus on one side of the visual field induced a motion aftereffect in the opposite visual field. The presence of a global motion aftereffect suggests that feature-based attention modulated the firing rates of direction-selective neurons with spatial receptive fields in the unstimulated region of the visual field, consistent with the data reported in Figures 2E and 2F.

The present report of feature-based attention spreading across the visual scene is also reminiscent of reports showing that mental imagery can selectively influence neural activity within retinotopically organized regions of visual cortex (e.g., Slotnick et al., 2005). On this account, “imagining” a stimulus might be accomplished by boosting the gain of neurons tuned to the imagined features. If so, then feature-based attention may be the mechanism of mental imagery, and future studies might investigate this link by using multivariate pattern classification methods to evaluate feature-based modulations within visual cortex during epochs of imagined stimulation.

Robust feature-selective attentional modulations were observed in ROIs that are not generally thought to play an important role in motion processing (e.g., V2v, V3v, V4v; see Figure 2). However, this observation is not unique to our study, as a recent report also demonstrated motion-selective modulations in fMRI response patterns from areas V1, V2, V3, V4, and hMT+ using fMRI and eight different directions of motion (Kamitani and Tong, 2006). In addition, robust classification accuracy need not be based solely on direction-selective neural activity per se. For example, some visual areas exhibit a biased population response in favor of radial motion as compared with circular motion (Beardsley and Vaina, 2005; Bex and Makous, 1997; Sasaki et al., 2006). It is also possible that although neurons in many of these regions are not thought to be particularly direction selective based on single-unit recording studies, the response profile across a population of such weakly selective neurons may in fact carry some information about direction. Thus, the multivariate fMRI response pattern may be sensitive to population-level dynamics that are not easily observed when recording the spiking activity of single neurons. Understanding the exact nature of these attentional modulations will be crucial for making claims about the tuning properties of neurons based on fMRI response patterns; however, our results support the general conclusion that a spatially global feature-based attentional mechanism gives rise to systematic changes in activation across visual cortex.

Feature-selective responses were also observed within IPS, even though this region is most often considered a “source” of attentional control signals instead of a site of modulation (Corbetta and Shulman, 2002). However, previous neuroimaging studies have demonstrated that IPS responds robustly to moving stimuli (Liu et al., 2003;

Shulman et al., 1999), and a recent single-unit study documented directional selectivity in monkey lateral IPS after training (Freedman and Assad, 2006). In conjunction with the present observations, these studies suggest that regions such as IPS (and perhaps FEF; see Figure 2A) may flexibly exhibit some feature selectivity, which in turn may facilitate the targeting of attentional modulations in earlier occipital visual areas. However, we endorse this explanation cautiously and future studies will need to confirm our results, as well as test the generality of feature-selective responses in parietal cortex.

Feature-based attentional modulations in the absence of direct visual stimulation might be mediated by a purely endogenous (or top-down) gain control mechanism, akin to “baseline shifts” in neural activity induced by space-based and feature-based attention shifts in the absence of a stimulus (Chawla et al., 1999; Hayden and Gallant, 2005; Kastner et al., 1999; Luck et al., 1997; Ress et al., 2000). However, previously reported baseline shifts were observed during the temporal gap between a cue instructing the observer where (or what) to attend and the presentation of the target stimulus or search array. In contrast, observers in the present experiment were continuously monitoring a stimulus on one side of the visual field, so the spread of feature-based attention may have been driven by hard-wired cross-hemispheric connections between similarly tuned neurons in corresponding visual areas. According to this model, similarly tuned neurons in each cortical hemisphere are connected in a mutually excitatory manner; the efficacy of these cross-hemispheric connections might be modified by attention, giving rise to feature-selective modulations in the absence of direct visual stimulation. In addition, inhibitory connections between corresponding visual areas in each hemisphere may play a role in producing feature-selective activation patterns. This possibility is consistent with the lack of generalization across both patterns of activation evoked by attended stimuli and patterns of activation measured in response to unstimulated regions of the visual field (see Results). While the exact nature of these connections remains unknown, it seems possible that they might be callosal, subcortical, or cortical in nature. Follow-up experiments examining the speed with which feature-based attention spreads across the visual field, as well as studies employing split-brain patients, may shed light on this issue.

According to another account, highly active neurons enhance the activity of corresponding neurons in complementary visual areas in the other hemisphere. Thus, the spread of feature-specific gain might depend only on the activity level of the “sending” neuron, which in the present study is determined by the priority of the attended feature. However, other factors that modulate the gain of the sending neuron, such as stimulus contrast, may also induce a global spread of feature-specific activity. In either case, the spread of feature-specific modulations likely plays a fundamental role in highlighting or priming relevant features across the visual field, and future studies will

need to determine the influence of various factors (e.g., attention, contrast, etc.) on this phenomenon.

While a better understanding of the exact mechanisms will clearly require additional study, the present results establish that feature-based modulations spread across the visual field, even to regions of space that do not contain a stimulus. The functional utility of this spatially global feature-based mechanism can easily be imagined: the ability to enhance sensitivity to a given feature across the visual field would facilitate efficient visual search and the detection of behaviorally relevant stimuli.

EXPERIMENTAL PROCEDURES

Participants

All participants gave written informed consent to participate in the study, which was approved by the Salk Institute Human Subjects Institute Review Board. Ten neurologically intact adults (four females), ages 25 to 30, participated in the main feature-based attention experiment. Four observers from the main experiment also participated in the spatial selectivity control experiment.

Feature-Based Attention Experiment

All visual stimuli were rendered in black on a white background and were viewed via the Avotec Silent Vision SV-701 Fiber Optic Visual System (Stuart, FL). Observers were instructed to maintain visual fixation on a central square that was present for the duration of each scan (subtending 0.16° visual angle). At the start of each trial, two small dots of the same size were presented 2.75° above fixation and ±5.9° to the left and right of fixation. An attention cue, consisting of a horizontal line subtending 0.4°, was present throughout each trial; the cue indicated both the location of the impending target stimulus and the attended direction of motion. The direction of the line indicated the location of the to-be-attended aperture. For half of the observers, a green central attention cue instructed observers to monitor the dots moving at 135° and a red cue instructed them to monitor the dots moving at 45°; this color/direction pairing was reversed for the remaining observers.

After 500 ms, moving dots were presented within either one or two invisible circular apertures (subtending 2.5° radius). Half of the dots in each aperture moved at 45° and the other half moved at 135° for the duration of the 14 s presentation period (each dot moved at 4.6°/s, subtended 0.2° diameter, and had a limited lifetime of eight 33.3 ms frames). When a single motion aperture was presented, it always appeared at the cued location. Target events were defined as a brief slowing (mean slowing of 2.9°/s) of the dots moving in the currently attended direction; distractor events were defined as a brief slowing of the dots moving in the unattended direction. Target and distractor events only occurred within the spatially attended aperture, and the speed of the dots in the unattended aperture (if present) remained constant throughout each trial. The observer's task was to press a button with the right pointer finger whenever a target event was detected. On each trial, there were two target and two distractor events, randomly interleaved and spaced 3.5 ± 0.2 s apart, with the first event starting 1 ± 0.2 s after the onset of the stimulus display. The next trial began after a blank intertrial interval of 6 s. There were 24 trials in each scanning run, and the pseudorandomized trial presentation order used for each scan was selected at random from a set of the 60 most statistically efficient sequences of events (based on 20,000 tested sequences; see Dale, 1999).

Independent Functional Localizer Task and Retinotopic Mapping

To identify spatially selective regions of occipital cortex, IPS, and FEF, a moving dot stimulus was presented in either the left or the right stimulus location for 16 s in an alternating sequence; the moving dot

stimulus contained either uncorrelated motion or a correlated motion flow field with the direction of motion changing every 2 s. Since spatially selective visual areas in IPS and FEF are known to be sensitive to attentional factors (Silver et al., 2005), observers were instructed to press a button when they detected a 1 s slowing in the motion of the dots, which occurred once per 16 s stimulation period. Retinotopic mapping data were obtained in 1–2 scans per observer using a checkerboard stimulus and standard presentation parameters (checkerboard flickering at 8 Hz and subtending 45° of polar angle; see Engel et al., 1994; Sereno et al., 1995). The data were projected onto computationally inflated cortical surfaces to aid in the visualization of early visual cortical areas revealed by the functional localizer and retinotopic mapping procedures.

Control Experiment: Spatial Selectivity of Visual ROIs

This control study was designed to assess the spatial selectivity of the voxels used in the main attention experiment within a subset of four observers. The size and structure of the stimuli and task exactly matched the functional localizer paradigm described above, except that epochs of peripheral stimulation were only 12 s in duration and were randomly interleaved with 12 s passive fixation trials in which no peripheral stimulus was presented. There were 20 stimulus-present trials and 10 passive fixation trials on each scan. The time-windowed average responses to contralateral and ipsilateral stimuli were computed by subtracting the mean response on passive fixation trials so that all time series reflect deviations away from the low-level fixation baseline (see also Figure S5 for a complementary analysis). Since the type of motion (uncorrelated or flow-field) only exerted a small main effect on activation levels in some regions (e.g., hMT+) and did not interact with spatial location, all data are presented collapsed across this factor.

fMRI Data Acquisition and Analysis

MRI scanning was performed on a Signa EXCITE 3 Tesla GE scanner equipped with an eight channel head coil at the Center for Functional Magnetic Resonance Imaging, University of California, San Diego. A custom-made bite bar was used to restrict head movement. Anatomical images were acquired using a SPGR T1-weighted sequence that yielded images with a 0.97 mm × 0.97 mm × 1 mm resolution. Whole-brain echoplanar functional images (EPI) were acquired in 32 transverse slices (TR = 2000 ms, TE = 30 ms, flip angle = 90°, image matrix = 64 × 64, FOV = 220 mm, slice thickness = 3 mm, no gap).

Data analysis was performed using BrainVoyager QX (v 1.74; Brain Innovation, Maastricht, The Netherlands) and custom time series analysis and pattern classification routines written in Matlab (v 7.1; The Math Works, Natick, MA). Data from the feature-based attention experiment were collected in either seven or eight scans per subject, with each scan lasting 486 s. EPI images were corrected using an unwarping procedure (the FUGE algorithm, FMRI Software Library, University of Oxford), slice-time corrected, motion-corrected (both within and between scans), and high-pass filtered (3 cycles/run) to remove low-frequency components in the time series. Data from the functional localizer task were collected in 1 or 2 scans. Each scan lasted 390 s and the EPI images were preprocessed as described above. The disparity in the number of scans for each observer was due only to constraints on scan time. In the spatial selectivity control experiment, observers participated in three 370 s scans.

ROI Selection

To identify ROIs in visual cortex that represented either the left or the right stimulus locations, a general linear model (GLM) with four boxcar regressors was applied to the BOLD time series data from the functional localizer scans; two regressors marked temporal epochs of correlated motion and uncorrelated motion on the right side of fixation, and two regressors marked temporal epochs of correlated and uncorrelated motion on the left side of fixation. Each of the boxcar regressors was then convolved with a gamma function to account for the

assumed hemodynamic response function ($\delta = 2.5$ s, $\tau = 1.25$ s; see Boynton et al., 1996). ROIs in V1, V2v, V3v, V4v, and V3A were defined by identifying voxels within each retinotopically defined area responding more strongly during epochs of visual stimulation on one side of space compared with the other (after collapsing across epochs of correlated and uncorrelated motion). hMT+ was defined as a contiguous group of voxels lateral to the parietal-occipital sulcus and beyond the retinotopically organized visual areas that exhibited a larger response during epochs of correlated motion compared with epochs of uncorrelated motion. ROIs in IPS were identified as contiguous clusters of spatially selective voxels superior to the parietal-occipital junction within the intraparietal sulcus (identified in 13/20 hemispheres), and ROIs in FEF were identified as contiguous clusters of the spatially selective voxels near the junction of the precentral sulcus and the superior frontal sulcus (identified in 11/20 hemispheres). The same ROIs were used in the feature-based attention experiment and in the control experiment.

In occipital cortex, the 50 most spatially selective voxels, as defined based on responses during the functional localizer task, were included in each ROI. By selecting the most selective voxels, the hMT+ ROIs were likely biased in favor of area MT, which is thought to be more spatially selective than area MST (Huk et al., 2002). However, we cannot rule out the possibility that the hMT+ ROIs also encompassed some portion of MST; hence, we adopt the more general hMT+ terminology. In IPS and FEF, we were not able to identify 50 activated voxels in all hemispheres, so all contiguous voxels that passed a minimum threshold of $p < 0.05$ (corrected for multiple comparisons using the false discovery rate method implemented in Brain Voyager) were included. The mean size (in voxels \pm SEM) of the IPS ROIs was 45.7 ± 2.0 , and the mean size of FEF ROIs was 40.3 ± 3.5 . Figure S7 shows the location of the IPS and FEF ROIs in each observer in native scanner space.

Pattern Classification Analysis

Our general image classification approach is similar to that reported elsewhere (see Haynes and Rees, 2005). We extracted seven EPI images on each trial (from 6–18 s after stimulus onset, where an image is defined as a vector of activation values from all voxels in a particular ROI). Before classification, the images extracted from all but one scan were defined as training images, and the remaining images were defined as test images. All test images belonging to a particular stimulus/attention configuration (see Figures 1B–1E) were then averaged together (after z-normalizing), creating a mean activation vector characterizing the pattern of responses across voxels for each condition in a scan. We then computed the md between the test activation vector (X) and each mean activation vector for the two attention conditions (attending 45° or 135° motion) computed over the training images:

$$md(i) = (X - \bar{X}_i)^T S^{-1} (X - \bar{X}_i)$$

where S is the pooled covariance matrix estimated from the training images, and \bar{X}_i is the mean training activation vector for each attention condition i . The parameters S and \bar{X}_i were computed using only data from the training set. The test activation vector was then assigned to the condition for which $md(i)$ was smallest. Classification was performed using activation vectors of different lengths, with the most discriminating voxels—determined using a pooled variance t test computed only using the training activation vectors (Haynes and Rees, 2005)—entered first (until all 50 voxels from the ROI were included). This procedure was repeated until each scan had served as a test set in turn.

Computing Classification Accuracy and Stimulus-Selectivity Index for Each Voxel in a Visual Area

To generate an estimate of classification accuracy for each voxel in a visual area, we defined a spherical neighborhood that encompassed

surrounding voxels (radius of the sphere = 2.5 voxels). Because we restricted our analysis to voxels in the cortical sheet, and because the cortical sheet is folded, the size of the neighborhood around each voxel was not identical (mean size in voxels, \pm SEM: 45 ± 14 , with neighborhoods of <10 voxels excluded). This sphere size was chosen to be as close to 50 voxels as possible so that the results would be roughly comparable to the data depicted in Figure 2. The results were also replicated using a sphere with a radius of 3 voxels (mean neighborhood size in voxels, \pm SEM: 60 ± 19 , data not shown). The classification accuracy for the voxel at the center of the sphere was computed using the same hold-one-scan out method described above. This procedure was iterated until a classification accuracy was assigned to each voxel within the visual area under consideration.

Estimates of stimulus selectivity for each voxel (Figure 4A) proceeded in an identical manner, except that instead of computing the classification accuracy for a neighborhood, we computed the magnitude of the response to a contralateral stimulus minus the response to an ipsilateral stimulus during the functional localizer scans (over a window extending from 6 s to 18 s poststimulus). In addition, for the four observers that participated in the control experiment where epochs of contralateral and ipsilateral stimulation were interleaved with fixation, we computed a selectivity index comparing contralateral stimulation versus fixation, providing a more direct measure of spatial selectivity (Figure S6). To produce the classification data depicted in Figure 4D, we sorted voxels into four bins based on a quartile analysis of the distribution of stimulus selectivity indices (shown in Figure 4C).

Group Random-Effects Analysis of Classification Accuracy in Occipital Cortex

For each voxel in the cortical sheet over the occipital lobe in each observer, we estimated classification accuracy using the method described in the preceding section (mean neighborhood size in voxels, \pm SEM: 45 ± 14). The classification accuracy map for each observer was then standardized into the space of Talairach and Tournoux (1988), and repeated-measures t tests were used to identify voxels that consistently classified the currently attended direction across observers. The single-voxel threshold was set at $t(9) = 2.3$, $p < 0.05$. Since this analysis was secondary to our main ROI-based analysis, and because the neighborhood analysis rendered the voxel-by-voxel t values nonindependent, we did not formally control for multiple comparisons. However, a minimum cluster size of 0.5 ml was adopted to partially guard against false positives (see, e.g., <http://afni.nimh.nih.gov/afni/doc/manual/AlphaSim>), and classification accuracy was only computed within the occipital cortical sheet (as opposed to the whole brain) because our main a priori theoretical interest was to evaluate the spread of feature-based attention in regions of early visual cortex. However, as an additional check to ensure that our algorithm did not always yield positive results, we repeated the analysis and estimated classification accuracy for all voxels in the image matrix—including voxels in white matter and outside of the brain. We found that the regions exhibiting above-chance classification accuracy were primarily confined to the cortical sheet (e.g., the areas reported in Figure 6 and Table 1), supporting the general validity of our analytical approach.

Eye Tracking

Eye tracking was performed at 60 Hz during scanning for three of the ten observers using an MR-compatible camera built into the video display goggles (Avotec Silent Vision SV-701 Fiber Optic Visual System). Data were first corrected for eye-blinks and linear drift, and then the mean position of the eye during each 14 s trial was separately assessed for each of the eight attentional conditions (Figure S4).

Supplemental Data

The Supplemental Data for this article can be found online at <http://www.neuron.org/cgi/content/full/55/2/301/DC1>.

ACKNOWLEDGMENTS

We thank John-Dylan Haynes and Lone Fine for helpful comments. This work was supported by NIH grants T32-MH20002, F32-EY017261, and EY12925. J.T.S. and G.M.B. designed research; J.T.S. performed research; J.T.S. analyzed data; and J.T.S. and G.M.B. wrote the paper.

Received: February 5, 2007

Revised: April 23, 2007

Accepted: June 5, 2007

Published: July 18, 2007

REFERENCES

- Albright, T.D., Desimone, R., and Gross, C.G. (1984). Columnar organization of directionally selective cells in visual area MT of the macaque. *J. Neurophysiol.* *51*, 16–31.
- Arman, A.C., Ciaramitaro, V.M., and Boynton, G.M. (2006). Effects of feature-based attention on the motion aftereffect at remote locations. *Vision Res.* *46*, 2968–2976.
- Awh, E., and Pashler, H. (2000). Evidence for split attentional foci. *J. Exp. Psychol. Hum. Percept. Perform.* *26*, 834–846.
- Beardsley, S.A., and Vaina, L.M. (2005). Psychophysical evidence for a radial motion bias in complex motion discrimination. *Vision Res.* *45*, 1569–1586.
- Ben Hamed, S., Duhamel, J.R., Bremmer, F., and Graf, W. (2001). Representation of the visual field in the lateral intraparietal area of macaque monkeys: a quantitative receptive field analysis. *Exp. Brain Res.* *140*, 127–144.
- Bex, P.J., and Makous, W. (1997). Radial motion looks faster. *Vision Res.* *37*, 3399–3405.
- Bichot, N.P., Rossi, A.F., and Desimone, R. (2005). Parallel and serial neural mechanisms for visual search in macaque area V4. *Science* *308*, 529–534.
- Boynton, G.M. (2005). Attention and visual perception. *Curr. Opin. Neurobiol.* *15*, 465–469.
- Boynton, G.M., Engel, S.A., Glover, G.H., and Heeger, D.J. (1996). Linear systems analysis of functional magnetic resonance imaging in human V1. *J. Neurosci.* *16*, 4207–4221.
- Braddick, O.J., O'Brien, J.M., Wattam-Bell, J., Atkinson, J., Hartley, T., and Turner, R. (2001). Brain areas sensitive to coherent visual motion. *Perception* *30*, 61–72.
- Chawla, D., Rees, G., and Friston, K.J. (1999). The physiological basis of attentional modulation in extrastriate visual areas. *Nat. Neurosci.* *2*, 671–676.
- Corbetta, M., and Shulman, G.L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nat. Rev. Neurosci.* *3*, 201–215.
- Cox, D.D., and Savoy, R.L. (2003). Functional magnetic resonance imaging (fMRI) “brain reading”: detecting and classifying distributed patterns of fMRI activity in human visual cortex. *Neuroimage* *19*, 261–270.
- Dale, A.M. (1999). Optimal experimental design for event-related fMRI. *Hum. Brain Mapp.* *8*, 109–114.
- Desimone, R., and Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annu. Rev. Neurosci.* *18*, 193–222.
- Desimone, R., and Ungerleider, L.G. (1986). Multiple visual areas in the caudal superior temporal sulcus of the macaque. *J. Comp. Neurol.* *248*, 164–189.
- Engel, S.A., Rumelhart, D.E., Wandell, B.A., Lee, A.T., Glover, G.H., Chichilnisky, E.J., and Shadlen, M.N. (1994). fMRI of human visual cortex. *Nature* *369*, 525.

- Folk, C.L., Leber, A.B., and Egeth, H.E. (2002). Made you blink! Contingent attentional capture produces a spatial blink. *Percept. Psychophys.* *64*, 741–753.
- Freedman, D.J., and Assad, J.A. (2006). Experience-dependent representation of visual categories in parietal cortex. *Nature* *443*, 85–88.
- Gandhi, S.P., Heeger, D.J., and Boynton, G.M. (1999). Spatial attention affects brain activity in human primary visual cortex. *Proc. Natl. Acad. Sci. USA* *96*, 3314–3319.
- Haenny, P.E., Maunsell, J.H., and Schiller, P.H. (1988). State dependent activity in monkey visual cortex. II. Retinal and extraretinal factors in V4. *Exp. Brain Res.* *69*, 245–259.
- Haxby, J.V., Gobbini, M.I., Furey, M.L., Ishai, A., Schouten, J.L., and Pietrini, P. (2001). Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science* *293*, 2425–2430.
- Hayden, B.Y., and Gallant, J.L. (2005). Time course of attention reveals different mechanisms for spatial and feature-based attention in area V4. *Neuron* *47*, 637–643.
- Haynes, J.D., and Rees, G. (2005). Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nat. Neurosci.* *8*, 686–691.
- Haynes, J.D., and Rees, G. (2006). Decoding mental states from brain activity in humans. *Nat. Rev. Neurosci.* *7*, 523–534.
- Huk, A.C., Dougherty, R.F., and Heeger, D.J. (2002). Retinotopy and functional subdivision of human areas MT and MST. *J. Neurosci.* *22*, 7195–7205.
- Kamitani, Y., and Tong, F. (2005). Decoding the visual and subjective contents of the human brain. *Nat. Neurosci.* *8*, 679–685.
- Kamitani, Y., and Tong, F. (2006). Decoding seen and attended motion directions from activity in the human visual cortex. *Curr. Biol.* *16*, 1096–1102.
- Kastner, S., De Weerd, P., Desimone, R., and Ungerleider, L.G. (1998). Mechanisms of directed attention in the human extrastriate cortex as revealed by functional MRI. *Science* *282*, 108–111.
- Kastner, S., Pinsk, M.A., De Weerd, P., Desimone, R., and Ungerleider, L.G. (1999). Increased activity in human visual cortex during directed attention in the absence of visual stimulation. *Neuron* *22*, 751–761.
- Liu, T., Slotnick, S.D., Serences, J.T., and Yantis, S. (2003). Cortical mechanisms of feature-based attentional control. *Cereb. Cortex* *13*, 1334–1343.
- Luck, S.J., Chelazzi, L., Hillyard, S.A., and Desimone, R. (1997). Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. *J. Neurophysiol.* *77*, 24–42.
- Macknik, S.L., and Livingstone, M.S. (1998). Neuronal correlates of visibility and invisibility in the primate visual system. *Nat. Neurosci.* *1*, 144–149.
- Martinez-Trujillo, J.C., and Treue, S. (2004). Feature-based attention increases the selectivity of population responses in primate visual cortex. *Curr. Biol.* *14*, 744–751.
- McMains, S.A., and Somers, D.C. (2004). Multiple spotlights of attentional selection in human visual cortex. *Neuron* *42*, 677–686.
- Melcher, D., Papatthomas, T.V., and Vidnyanszky, Z. (2005). Implicit attentional selection of bound visual features. *Neuron* *46*, 723–729.
- Mitchell, T.M., Hutchinson, R.S., Pereira, F., Wang, X., Just, M., and Newman, S. (2003). Learning to decode cognitive states from brain images. *Mach. Learn.* *57*, 145–175.
- Moran, J., and Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. *Science* *229*, 782–784.
- Motter, B.C. (1994). Neural correlates of feature selective memory and pop-out in extrastriate area V4. *J. Neurosci.* *14*, 2190–2199.
- Norman, K.A., Polyn, S.M., Detre, G.J., and Haxby, J.V. (2006). Beyond mind-reading: multi-voxel pattern analysis of fMRI data. *Trends Cogn. Sci.* *10*, 424–430.
- Orban, G.A., Fize, D., Peuskens, H., Denys, K., Nelissen, K., Sunaert, S., Todd, J., and Vanduffel, W. (2003). Similarities and differences in motion processing between the human and macaque brain: evidence from fMRI. *Neuropsychologia* *41*, 1757–1768.
- Peelen, M.V., and Downing, P.E. (2007). Using multi-voxel pattern analysis of fMRI data to interpret overlapping functional activations. *Trends Cogn. Sci.* *11*, 4–5.
- Ress, D., Backus, B.T., and Heeger, D.J. (2000). Activity in primary visual cortex predicts performance in a visual detection task. *Nat. Neurosci.* *3*, 940–945.
- Saenz, M., Buracas, G.T., and Boynton, G.M. (2002). Global effects of feature-based attention in human visual cortex. *Nat. Neurosci.* *5*, 631–632.
- Saenz, M., Buracas, G.T., and Boynton, G.M. (2003). Global feature-based attention for motion and color. *Vision Res.* *43*, 629–637.
- Sasaki, Y., Rajimehr, R., Kim, B.W., Ekstrom, L.B., Vanduffel, W., and Tootell, R.B. (2006). The radial bias: a different slant on visual orientation sensitivity in human and non human primates. *Neuron* *51*, 661–670.
- Serences, J.T., and Yantis, S. (2006). Selective visual attention and perceptual coherence. *Trends Cogn. Sci.* *10*, 38–45.
- Serences, J.T., Shomstein, S., Leber, A.B., Golay, X., Egeth, H.E., and Yantis, S. (2005). Coordination of voluntary and stimulus-driven attentional control in human cortex. *Psychol. Sci.* *16*, 114–122.
- Sereno, M.I., Dale, A.M., Reppas, J.B., Kwong, K.K., Belliveau, J.W., Brady, T.J., Rosen, B.R., and Tootell, R.B. (1995). Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Science* *268*, 889–893.
- Shulman, G.L., Ollinger, J.M., Akbudak, E., Conturo, T.E., Snyder, A.Z., Petersen, S.E., and Corbetta, M. (1999). Areas involved in encoding and applying directional expectations to moving objects. *J. Neurosci.* *19*, 9480–9496.
- Silver, M.A., Ress, D., and Heeger, D.J. (2005). Topographic maps of visual spatial attention in human parietal cortex. *J. Neurophysiol.* *94*, 1358–1371.
- Slotnick, S.D., Thompson, W.L., and Kosslyn, S.M. (2005). Visual mental imagery induces retinotopically organized activation of early visual areas. *Cereb. Cortex* *15*, 1570–1583.
- Talairach, J., and Tournoux, P. (1988). *Co-planar Stereotaxic Atlas of the Human Brain* (New York: Thieme).
- Tootell, R.B., Mendola, J.D., Hadjikhani, N.K., Ledden, P.J., Liu, A.K., Reppas, J.B., Sereno, M.I., and Dale, A.M. (1997). Functional analysis of V3A and related areas in human visual cortex. *J. Neurosci.* *17*, 7060–7078.
- Tootell, R.B., Mendola, J.D., Hadjikhani, N.K., Liu, A.K., and Dale, A.M. (1998). The representation of the ipsilateral visual field in human cerebral cortex. *Proc. Natl. Acad. Sci. USA* *95*, 818–824.
- Treue, S., and Martinez Trujillo, J.C. (1999). Feature-based attention influences motion processing gain in macaque visual cortex. *Nature* *399*, 575–579.
- Treue, S., and Maunsell, J.H. (1996). Attentional modulation of visual motion processing in cortical areas MT and MST. *Nature* *382*, 539–541.
- Williford, T., and Maunsell, J.H. (2006). Effects of spatial attention on contrast response functions in macaque area V4. *J. Neurophysiol.* *96*, 40–54.
- Wolfe, J.M. (1994). Guided Search 2.0: A revised model of visual search. *Psychon. Bull. Rev.* *1*, 202–238.