

A Hierarchy of Attentional Priority Signals in Human Frontoparietal Cortex

Taosheng Liu^{1,2} and Youyang Hou¹

¹Department of Psychology and ²Neuroscience Program, Michigan State University, East Lansing, Michigan 48824

Humans can voluntarily attend to a variety of visual attributes to serve behavioral goals. Voluntary attention is believed to be controlled by a network of dorsal frontoparietal areas. However, it is unknown how neural signals representing behavioral relevance (attentional priority) for different attributes are organized in this network. Computational studies have suggested that a hierarchical organization reflecting the similarity structure of the task demands provides an efficient and flexible neural representation. Here we examined the structure of attentional priority using functional magnetic resonance imaging. Participants were cued to attend to location, color, or motion direction within the same stimulus. We found a hierarchical structure emerging in frontoparietal areas, such that multivoxel patterns for attending to spatial locations were most distinct from those for attending to features, and the latter were further clustered into different dimensions (color vs motion). These results provide novel evidence for the organization of the attentional control signals at the level of distributed neural activity. The hierarchical organization provides a computationally efficient scheme to support flexible top-down control.

Introduction

Goal-directed behavior requires the selective processing of task-relevant information in complex environments. Visual selective attention allows us to focus on specific aspects of the scene for prioritized processing (Reynolds and Chelazzi, 2004; Carrasco, 2011). Contemporary theories of attention have assumed a role of top-down control signals that bias bottom-up sensory processing (Wolfe, 1994; Desimone and Duncan, 1995). Neuroimaging studies have implicated a network of brain areas in dorsal frontoparietal cortex during top-down attentional control (Kastner and Ungerleider, 2000; Corbetta and Shulman, 2002). A possible function of these areas is to maintain attentional priority, i.e., the behavioral importance of items. In the case of spatial attention (selection of spatial locations), the idea has been strongly supported by data from both single-unit physiology and neuroimaging (Moore, 2006; Serences and Yantis, 2006; Bisley and Goldberg, 2010). More recent studies using the functional magnetic resonance imaging (fMRI) multivariate decoding approach have suggested that these areas also represent attentional priority for features and objects (Liu et al., 2011; Guo et al., 2012; Hou and Liu, 2012). These findings suggest that the dorsal frontoparietal areas represent attentional priority for locations, features, and objects.

An important question arises concerning the relationship between different types of priority signals within the same network. In particular, the relationship between priority signals for spatial and nonspatial properties has not been elucidated. In terms of the modulatory effect of attention, it has been shown that spatial and feature-based attention have distinct influences on both psychophysical performance (Baldassi and Verghese, 2005; Liu et al., 2007; Ling et al., 2009) and sensory responses (Saenz et al., 2002; Martinez-Trujillo and Treue, 2004). These considerations thus pose a challenge in understanding attentional control, that is, how a domain-general control mechanism provides domain-specific modulations.

Theoretical and computational studies have provided two broad perspectives on the organization of neural representations. On the one hand, it has been suggested that similar objects/concepts are represented by similar neural substrates, which could form a hierarchical structure for certain domains, such as object categories and action sequences (Edelman, 1998; Cooper and Shallice, 2006). On the other hand, nonhierarchical representations have also been proposed to account for complex behavior (Botvinick and Plaut, 2004). Although these studies are not concerned with attentional control per se, the differing views imply different organizational schemes for attentional priority. The former could predict a hierarchical organization such that priority signals for features within a dimension are more similar to each other than across dimensions, and at a higher level, priority signals for features in general are more similar to each other than their similarity to priority signals for locations. Alternatively, if neural representations are nonhierarchical, this would predict a lack of systematic organization for priority signals.

We hypothesized that priority signals for different properties are organized in a hierarchical fashion, given such an organization can provide both flexible and specific representations (Hin-

Received April 26, 2013; revised Aug. 29, 2013; accepted Sept. 10, 2013.

Author contributions: T.L. and Y.H. designed research; Y.H. performed research; T.L. and Y.H. analyzed data; T.L. wrote the paper.

This work was supported in part by a National Institutes of Health Grant (R01EY022727). We thank the Department of Radiology at Michigan State University for support of brain imaging research, and Ms. Scarlett Doyle and Dr. David Zhu for assistance in data collection. We also thank Dr. Pang-Ning Tan for advice on data analysis, and Drs. Justin Gardner and Jeremy Gray for helpful comments on this manuscript.

Correspondence should be addressed to Taosheng Liu, Department of Psychology, Michigan State University, 316 Physics Road, East Lansing, MI 48824. E-mail: tsliu@msu.edu.

DOI:10.1523/JNEUROSCI.1780-13.2013

Copyright © 2013 the authors 0270-6474/13/3316606-11\$15.00/0

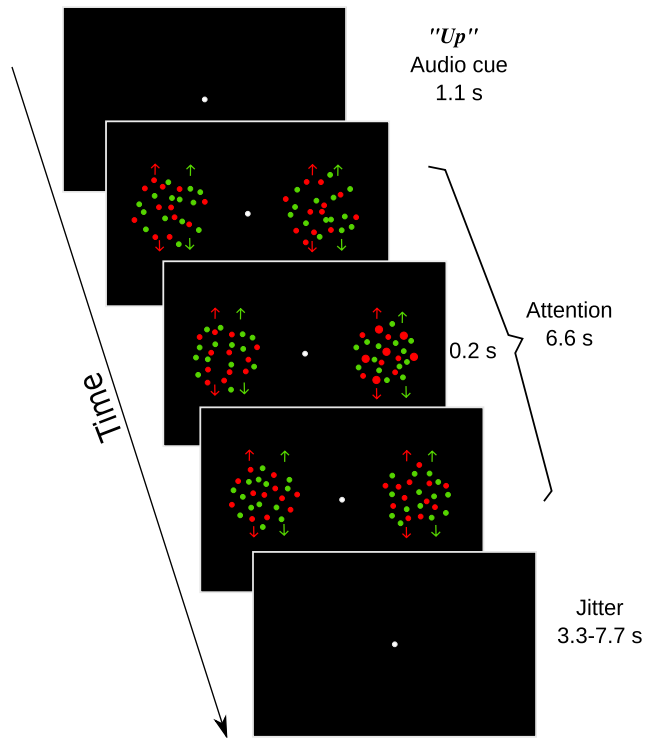


Figure 1. Schematic of an “up” trial in the attention task. Arrows denote the moving direction of dots, which were not shown in the actual stimuli. The third section illustrates a target event, in which the upward-moving red dots in the right aperture had a transient change (0.2 s) in dot size.

ton et al., 1986; Edelman, 1998). We manipulated top-down attention to spatial locations and visual features in a single experiment and analyzed multivoxel pattern similarity associated with different types of attention. The multivariate technique represents an information-based approach in analysis of neural data, which is critical in understanding the structure of the underlying neural representations (Kriegeskorte et al., 2006).

Materials and Methods

Participants

Twelve individuals (6 females) participated in the experiment; all had normal or corrected-to-normal vision; 11 were right-handed and 1 was left-handed. Two of the participants were authors, the rest were graduate and undergraduate students at Michigan State University. Participants were paid for their participation and gave informed consent according to the study protocol, which was approved by the Institutional Review Board at Michigan State University.

Stimulus and display

The visual stimuli consisted of two circular apertures (9° in diameter), each containing coherently moving dots (dot size: 0.18°), centered 8° to the left and right of a central fixation point (0.3° diameter) on a black background. In each of the two apertures, half of the dots were rendered in red and the other half in green; within each color group, half of the dots moved upward and the other half moved downward. Thus there were eight dot groups in total, generated by the combination of 2 locations \times 2 colors \times 2 directions, with each group of dots containing 15 dots. The speed of dot groups varied between 1.7 and $2.5^\circ/\text{s}$, with speed randomly assigned to each dot group on each trial (Fig. 1).

All stimuli were generated using MGL (<http://gru.brain.riken.jp/doku.php?id=mgl:overview>), a set of custom OpenGL libraries running in MATLAB (MathWorks). Images were projected on a rear-projection screen located in the scanner bore by a Toshiba TDP-TW100U projector outfitted with a custom zoom lens (Navitar). The screen resolution was

set to 1024×768 and the display was updated at 60 Hz. Participants viewed the screen via an angled mirror attached to the head coil at a viewing distance of 60 cm. Auditory stimuli were delivered via headphones by a Serene Sound Audio System (Resonance Technology).

Task and design

Attention experiment. Participants were instructed to fixate on the central disk throughout the experiment. At the beginning of each trial, an audio cue was played through the headphones worn by participants. There were three types of cues: two location cues (“left” or “right”) instructed participants to maintain attention on dots in either the left or right aperture, regardless of their color or direction, two color cues (“red,” “green”) instructed participants to maintain attention on either the red or green dots, regardless of their location or direction, and two direction cues (“up” or “down”) instructed participants to maintain attention to either upward-moving dots or downward-moving dots, regardless of their location or color. Thus, with each attentional cue, participants need to attend to four dot groups, out of eight possible dot groups. The audio cues were prerecorded with a native English speaker and stored as digital files on the computer. At 1.1 s after the onset of the audio cue, the dot stimuli appeared for 6.6 s. At random times during this interval, all dots in one or two of the dot groups briefly increased their size (Fig. 1). The size change either occurred on one of the four cued dot groups (target) or one of the four uncued dot groups (distracter). Participants were instructed to press a button when they detected a target, and to withhold response for distracters. For example, in the “up” trials, participants needed to attend to all the upward-moving dots and press the button when they noticed any upward-moving dots increased their size. On each trial, there was either one target only (a size-change event in one of the cued dot groups), one distracter only (a size-change event in one of the uncued dot groups), or one target and one distracter (both aforementioned size-change events). We counted a key press within a 1.5 s window after a target or distracter as a positive response to that event (on trials containing both a target and a distracter, the two events were separated by at least 1.5 s). A jittered intertrial interval followed the dot stimuli (3.3–7.7 s). In each scanning run, there were 4 trials for each cue condition, for a total of 24 trials. Trial order was randomly determined for each run. Participants performed 10 runs in the scanner, resulting in a total of 240 trials, with 40 trials per cue condition.

At the beginning of each scanning session, the red and green colors were set at isoluminance via heterochromatic flicker photometry (Kaiser, 1991). We fixed the red color to the RGB value of [255 0 0] and let participants adjust the green value to achieve minimum flicker. Their individual color setting was then used for the dot stimuli in the attention experiment. We also ran a threshold task to determine the appropriate magnitude of the size change for the dot stimuli. The task was identical to the attention task described above, except that the magnitude of size increment was controlled via three separate 1-up 2-down staircases, one for each attention type (location, color, motion). We fitted the staircase data with Weibull functions and selected size increments that yielded $\sim 80\%$ correct performance for each of the three attention types.

Practice and eye tracking. Each participant practiced the attention task in the behavioral lab for at least 1.5 h before the fMRI scan. The first part of practice consisted of setting isoluminant red/green colors and the size change threshold task as described above. Once participants achieved stable thresholds over several runs, we fixed the size changes and practiced them in the scanner version of the task. During these practice trials, we also monitored their eye position with an EyeLink II system (SR Research). All participants took part in the eye tracking session, with each performing at least two runs of the attention task. If we observed any subtle differences in fixation patterns across conditions, we gave further instructions and let participants practice more runs, until no systematic difference could be observed.

Retinotopic mapping. Early visual cortex and posterior parietal areas containing topographic maps were defined in a separate scanning session for each participant. We used rotating wedge and expanding/contracting rings to map the polar angle and radial component, respectively (Sereno et al., 1995; DeYoe et al., 1996; Engel et al., 1997). Borders between visual areas were defined as phase reversals in a polar angle map of the visual

field. Phase maps were visualized on computationally flattened representations of the cortical surface, which were generated from the high-resolution anatomical image using FreeSurfer and custom MATLAB code. In addition to occipital visual areas, our retinotopic mapping procedure also identified topographic areas in the intraparietal sulcus (IPS), IPS 1–4 (Swisher et al., 2007). In a separate run, we also presented moving versus stationary dots in alternating blocks and localized the human motion-sensitive area, hMT+, as an area near the junction of the occipital and temporal cortex that responded more to moving than stationary dots (Watson et al., 1993). Thus for each participant, we identified the following areas: V1, V2, V3, V3AB, V4, V7, hMT+, IPS1, IPS2, IPS3, and IPS4. There is controversy regarding the definition of visual area V4 (for review, see Wandell et al., 2007). Our definition of V4 followed that of Brewer et al. (2005), which defines V4 as a hemifield representation directly anterior to V3v.

MRI data acquisition

All functional and structural brain images were acquired using a GE Healthcare 3.0T Signa HDx MRI scanner with an 8-channel head coil, in the Department of Radiology at Michigan State University. For each participant, high-resolution anatomical images were acquired using a T1-weighted MP-RAGE sequence (FOV = 256 mm × 256 mm, 180 sagittal slices, 1 mm isotropic voxels) for surface reconstruction and alignment purposes. Functional images were acquired using a T2*-weighted echo planar imaging sequence consisted of 30 slices (TR = 2.2 s, TE = 30 ms, flip angle = 80°, matrix size = 64 × 64, in-plane resolution = 3 mm × 3 mm, slice thickness = 4 mm, interleaved, no gap). In each scanning session, a 2D T1-weighted anatomical image was also acquired that had the same slice prescription as the functional scans, for the purpose of aligning functional data to high resolution structural data.

fMRI data analysis

Data were processed and analyzed using mrTools (<http://www.cns.nyu.edu/heegerlab/wiki/doku.php?id=mrtools:top>) and custom code in MATLAB. Preprocessing of function data included head movement correction, linear detrend, and temporal high-pass filtering at 0.01 Hz. The functional images were then aligned to high-resolution anatomical images for each participant. Functional data were converted to percentage signal change by dividing the time course of each voxel by its mean signal over a run, and data from the 10 scanning runs were concatenated for subsequent analysis. All regions of interest (ROI) analyses were performed on individual participant's native anatomical space. For group-level analysis, we used surface-based spherical registration as implemented in Caret to coregister the individual participant's functional data to the Population-Average, Landmark- and Surface-based (PALS) atlas (Van Essen, 2005). Group-level statistics (random effects) were computed in the atlas space and the statistical parameter maps were visualized on a standard atlas surface (the "very inflated" surface). To correct for multiple comparisons, we set the threshold of the maps based on individual voxel level p value in combination with a cluster constraint, using the 3dClustSim program distributed as part of AFNI (http://afni.nimh.nih.gov/pub/dist/doc/program_help/3dClustSim.html). More details of the group analysis pipeline have been reported in previous studies (Liu et al., 2011; Hou and Liu, 2012).

Univariate analysis: deconvolution. Each voxel's time series was fitted with a general linear model whose regressors contained six attention conditions (left, right, red, green, up, and down). Each regressor modeled the fMRI response in a 25 s window after trial onset. The design matrix was pseudo-inversed and multiplied by the time series to obtain an estimate of the hemodynamic response for each attention condition. To measure the response magnitude of a region, we averaged the deconvolved response across all the voxels in an ROI.

In addition to the visual and parietal regions defined by retinotopic mapping, we also defined ROIs active during the attention task. This was done by using the goodness of fit measure (r^2 value), which is the amount of variance in the fMRI time series explained by the deconvolution model. The statistical significance of the r^2 value was evaluated via a permutation test by randomizing event times and recalculating the r^2 value using the deconvolution model (Gardner et al., 2005). One thou-

sand permutations were performed and the largest r^2 value in each permutation formed a null distribution expected at chance (Nichols and Holmes, 2002). Each voxel's p value was then calculated as the percentile of voxels in the null distribution that exceeded the observed r^2 value of that voxel. Using a cutoff p value of 0.05, we defined three additional areas that were active during the attention task: auditory cortex (AUD), frontal eye field (FEF), inferior frontal junction (IFJ) in both hemispheres. We were able to define these three areas on each individual participant's map. Additional idiosyncratic activations in individual maps were not further investigated as they were not consistently active across participants and hence absent in the group map (see below).

For group analysis, we transformed the individually obtained r^2 maps to the PALS atlas space and averaged their values (see Fig. 3). The null distributions of each individual participant were combined and the aggregated distribution served as the null distribution to obtain a voxelwise p value associated with an averaged r^2 value. We used a p of 0.02 and a cluster extent of 15 to threshold the group r^2 map, which corresponded to a whole-brain false positive rate of 0.01 according to 3dClustSim. Note this group average analysis was for visualization purpose only.

Univariate analysis: whole-brain contrast. To localize cortical areas differentially activated by different attention cues, we performed three linear contrasts analysis (location vs color, location vs direction, and color vs direction). We first transformed the individually estimated hemodynamic response to the atlas space, and then performed t tests comparing the average response magnitude for time points 3–6 (6.6–13.2 s), which spanned the peak time of the hemodynamic response (see Fig. 4). We combined the two cues within each attention type, e.g., for location–color comparison, the contrast vector was left + right–(red + green). We also compared individual cues within dimensions to examine differential brain activity for specific cues. Three linear contrasts were performed: left versus right, red versus green, up versus down. The resulting t -maps were thresholded at the same level as the r^2 map above (voxelwise $p = 0.02$, cluster extent = 15). This analysis revealed brain areas differentially activated by different attention cues (see Fig. 5). For illustration purposes only, we also obtained the grand average of hemodynamics responses (across voxels and participants) in select areas to show the relationship between different conditions.

Multivariate analysis: multivoxel similarity and cluster analyses. All multivariate analyses were performed in individually defined ROIs, either via retinotopic mapping or task-related activation using the r^2 criterion (see above). The retinotopically defined ROIs were further restricted to voxels whose r^2 values exceeded 0.06. In practice this thresholding procedure excluded most noisy voxels not stimulated by the visual stimuli (e.g., part of V1 corresponding to the far periphery), while being lenient enough to retain a sufficient number of voxels for multivariate pattern analyses. For each voxel and each attention condition, we first averaged time points 3–6 in their fMRI response to obtain a response amplitude. We then averaged this response amplitude across voxels in that ROI and subtracted the average response from each voxel's original response. This resulted in a list of values, with each number being the normalized response amplitude of a voxel in that condition, which we refer to as the response vector. For each participant and each ROI, we thus obtained six response vectors, one for each attention condition. The normalization procedure ensured that the analysis was insensitive to mean differences across different conditions and/or ROIs. Multivoxel similarity was measured by computing the correlation coefficients between all possible pairs of these response vectors (a total of 15 pairings). Note we could also use pairwise classification accuracy as a measure of pattern similarity, which indeed yielded largely similar results as reported here. However, simple correlation offered a more direct estimate of similarity, as classification accuracy depends on other analytical factors (e.g., the choice of classifiers and their parameters).

To interpret the correlation results, it is also necessary to have some measure of the reliability (i.e., the stability of the activity pattern for the same condition). Reliability was computed using the Spearman–Brown formula (Nunnally, 1978), which assessed the degree of self-correlation given the noise in the data. Specifically, a split-half reliability was calculated by correlating the response vector from a random half of the data with that from the other half. This was repeated 50 times, and the average

correlation (r) was used in the following formula: $r' = 2r/(1+r)$, to obtain the corrected reliability (r').

After obtaining the similarity measures for each ROI in each individual participant, we averaged the similarity measures across individuals. The resulting similarity measures were then used to compute a distance measure, defined as $1 - r$. Finally, clustering analyses were performed on the distance measures, to assign conditions into groups (clusters) such that conditions within groups are more similar to each other than conditions in different groups. We used the complete linkage algorithm to build the cluster structure, which used the largest distance between objects to separate two clusters (Tan et al., 2005). The results of clustering analyses were visualized by plotting dendrograms, which can reveal any hierarchical structure among different conditions. The stability of these structures was assessed by a permutation test. In the first test, we randomly sampled half of the trials for each participant and performed the multivoxel similarity analysis. We repeated this procedure for 50 times and averaged the similarity values upon which we performed the same clustering analysis as described above. In the second test, we again randomly sampled half of the trials, but with the trial labels randomly shuffled among conditions. This was repeated 50 times and we again obtained average similarity values and its associated clustering results.

To further assess the statistical significance of the cluster structures, we divided the similarity matrix into different partitions suggested by the clustering results (see Fig. 9). We then averaged the correlation values in those partitions for each participant and applied Fisher's z -transform to convert them into normally distributed z values. Finally, we compared the average correlations between partitions via paired t tests to evaluate the statistical significance.

Results

Behavior

Behavioral data were analyzed using both attention type (location, color, motion) and cue condition (left, right, red, green, upward, downward) as factors. We first compared the magnitude of the size change for different attention types, whose values were determined by separate staircases (see Materials and Methods). We did not find any significant differences in size-change magnitude ($F_{(2,22)} = 2.66, p > 0.09$). For the change detection task, we found equivalent performance across attention types (Fig. 2); this was expected given we thresholded the detection task at similar performance level. For accuracy scores (Hit–False alarm), there was no significant effect of attention type ($F_{(2,22)} < 1$) or cue condition ($F_{(5,55)} = 1.21, p > 0.31$). For reaction time (RT), there was no significant effect of attention type ($F_{(2,22)} = 1.53, p > 0.20$), but there was a significant effect of cue condition ($F_{(5,55)} = 3.21, p < 0.05$). Pairwise comparisons showed that reaction time for the green cue condition (842 ± 112 ms, mean \pm SD) was significantly faster than those for the right (906 ± 136 ms), red (935 ± 115 ms), up (905 ± 99 ms), and down (937 ± 104 ms) cue conditions (all p s < 0.05), with no other comparisons reaching significance. Note, however, that this difference in RT did not correspond to differences in fMRI multivariate patterns reported below. These results showed that participants were able to attend to the cued group of dots and ignore the uncued group of dots, and that task difficulty was similar across conditions.

It has been pointed out recently that averaging results from multivariate pattern analysis might be confounded by individual-level behavioral variations across conditions, that, although exhibiting inconsistent patterns across participants, could nevertheless produce significant group-level effects because the statistics are directionless (Todd et al., 2013). To address this potential confound, we examined individual participant's behavior data. Because our multivariate analyses showed clustering based on attention type (see results below), we focused our individual analyses on potential differences across the three attention types.

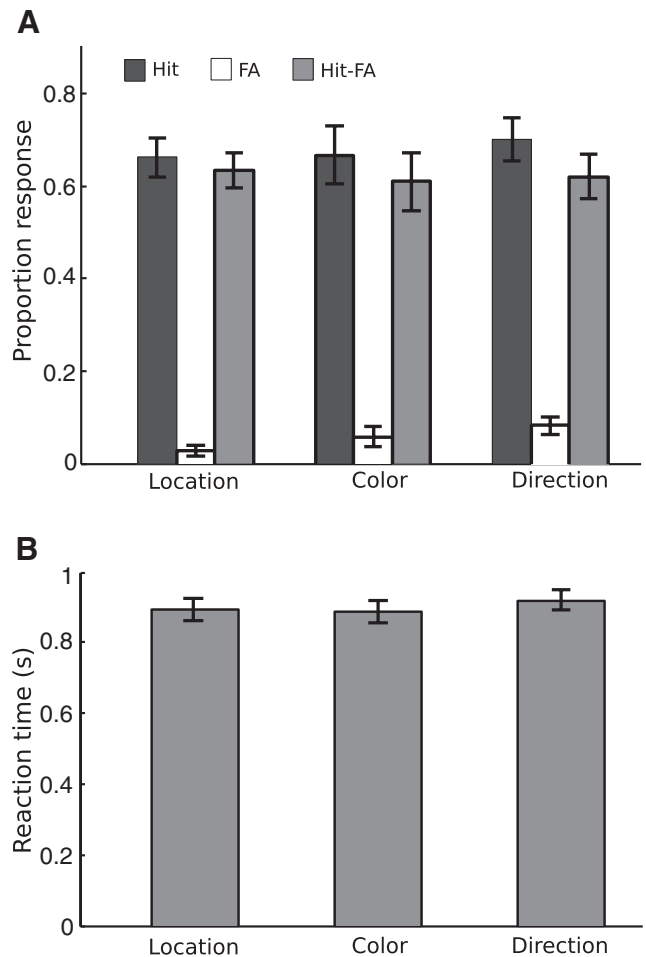


Figure 2. Behavioral results in the scanner. **A**, Mean proportion of response for hit, false alarm (FA), and hit–false alarm. **B**, Mean reaction time. Error bars indicate ± 1 SEM across participants ($N = 12$).

For each participant, we obtained a 2×3 contingency table for the number of hits and misses in the three conditions. We then performed a χ^2 test for independence and found 4 of 12 participants showed a significant effect ($p < 0.05$). In two participants, color cues led to fewer hits than the other cues; in one participant, color cues led to more hits than the other cues; and in one participant, location cues led to fewer hits than the other cues. We did not perform this test for false alarm and correct rejection trials, as false alarms were quite rare (many cells had zero frequency). We also analyzed RT for each participant, treating RT on single trials as input data to a one-way ANOVA. This analysis yielded a significant effect in 3 of 12 participants ($p < 0.05$). Paired comparison showed that one participant responded faster for location than direction cues, one participant showed the opposite pattern, and another participant responded faster for color than direction cues. Overall, behavioral differences across conditions were only present in a minority of our participants, and these differences did not correspond to the hierarchical structure we found in the fMRI data (see results below). Finally, when we removed participants who showed a significant effect in either accuracy or RT (6 of 12 participants), we obtained essentially the same group-level fMRI results in frontoparietal areas. We thus conclude that behavioral variations at individual-subject level cannot account for the results from fMRI multivariate analyses.

To verify participants maintained central fixation in our task, we monitored their eye position during practice sessions in the behavioral lab. Eye position was averaged across trials for each participant and subjected to ANOVA, which revealed no significant difference among the three attention types (location, color, motion), for either the horizontal ($F_{(2,22)} < 1$) or vertical ($F_{(2,22)} = 1.34, p > 0.28$) eye position. When we compared eye position among the six cue conditions, there was again no difference in either horizontal ($F_{(5,55)} = 1.05, p > 0.39$) or vertical ($F_{(5,55)} = 1.76, p > 0.14$) eye position. Thus, participants were able to maintain their fixation during the experiment and there was no systematic difference between fixation behaviors for different attention conditions.

Overall brain activity during the task

We first defined cortical areas whose activities were modulated by the attention task, using the goodness of fit criterion r^2 (see Materials and Methods). To visualize these areas, we projected the average r^2 map to the atlas surface (Fig. 3), which showed activity in a network of areas in occipital, parietal, and frontal cortex. The occipital activity coincided with retinotopically defined visual areas (V1, V2, V3, V3AB, V4, V7, and hMT+), and the parietal activity ran along the IPS, which coincided with retinotopically defined IPS areas (IPS1–4). Because individual IPS areas are quite small, for all following analyses we combined the four IPS areas into two areas IPS12 and IPS34, to increase power and to simplify data presentation. Frontal activity included a region around posterior superior frontal sulcus and precentral sulcus, the putative human FEF (Paus, 1996) and an area in the posterior inferior frontal sulcus and precentral sulcus, which we refer to as the IFJ. In addition, an area in the superior temporal cortex was also active, which was consistent with the human AUD. All these areas were found in both hemispheres, in a bilaterally symmetric pattern. Thus, for each participant, we obtained retinotopically defined occipital and parietal areas, as well as three task-defined areas: FEF, IFJ, and AUD. Note the r^2 method is an unbiased criterion to select voxels whose activities were consistently modulated by the task, regardless of their relative response amplitude between conditions. To increase power and simplify data presentation, we combined the corresponding ROIs in the left and right hemisphere in the subsequent analyses (unless otherwise noted). ROIs were combined at the level of functional data, not at the level of summary statistics.

We next examined the time courses in individually defined ROIs. For this analysis, we averaged fMRI response for each attention type (location, color, motion) across all voxels in an ROI. All areas showed an increase in fMRI response relative to the baseline (fixation during intertrial interval). Figure 4 shows fMRI time course from eight select ROIs. We compared the average response amplitude among three attention types (location, color,

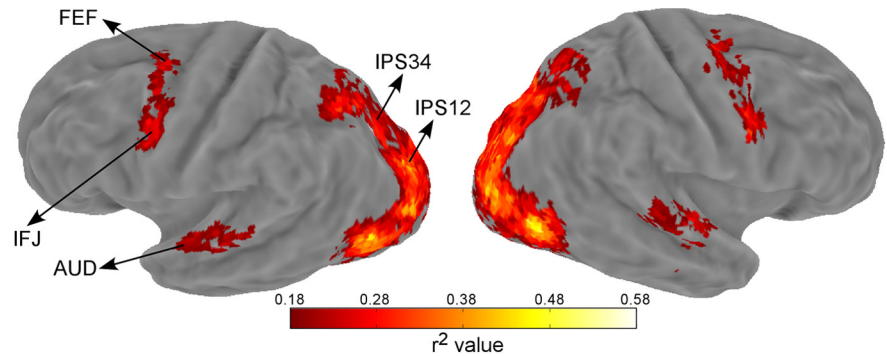


Figure 3. Group-averaged r^2 map shown on an inflated Caret atlas surface. The approximate locations of the three task-defined areas (AUD, FEF, IFJ) and two combined IPS regions (IPS12, IPS34) were indicated by arrows. AUD: auditory cortex, FEF: frontal eye field, IFJ: inferior frontal junction.

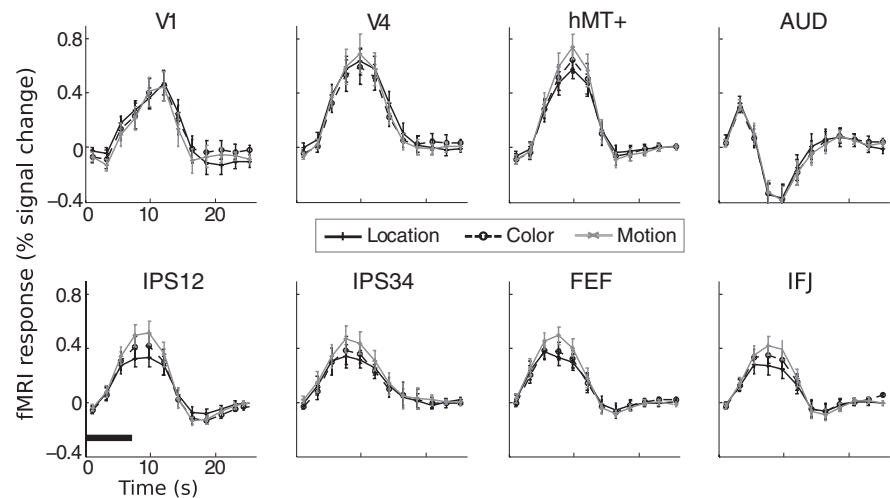


Figure 4. Mean time course data from eight select regions of interest. Error bars indicate ± 1 SEM across participants. The horizontal bar in the lower left section indicates the duration of the visual stimuli on a trial.

direction) using one-way repeated-measures ANOVA. We observed a significant effect in six ROIs: V7 ($F_{(2,22)} = 3.56, p < 0.05$), hMT+ ($F_{(2,22)} = 4.21, p < 0.05$), IPS12 ($F_{(2,22)} = 7.24, p < 0.01$), IPS34 ($F_{(2,22)} = 5.89, p < 0.01$), FEF ($F_{(2,22)} = 5.36, p < 0.05$), and IFJ ($F_{(2,22)} = 7.71, p < 0.01$). In those areas, attention to motion evoked a larger response than attention to color and location. In the remaining ROIs, the three attention types elicited equivalent levels of fMRI response.

Because spatial attention is known to retinotopically modulate cortical activity, we determined whether the location cue in our task produced a similar effect. For this analysis, we examined time course for the “attend to left” and “attend to right” trials separately for the left and right hemisphere ROIs. In general, we found a contralateral attentional modulation: attending to the left aperture produced a larger fMRI response than attending to the right aperture in right hemisphere ROIs, and vice versa (data not shown). We performed a two-way repeated-measures ANOVA on average peak amplitude with factors of spatial attention (attend to left vs attend to right) and hemisphere (left vs right). This analysis revealed significant interaction in V3, V3AB, IPS12, FEF, and IFJ (all $p < 0.05$). These results are consistent with previous research showing the contralateral modulation due to spatial attention (Kastner and Ungerleider, 2000; Carrasco, 2006), indicating our manipulation of spatial attention was effective.

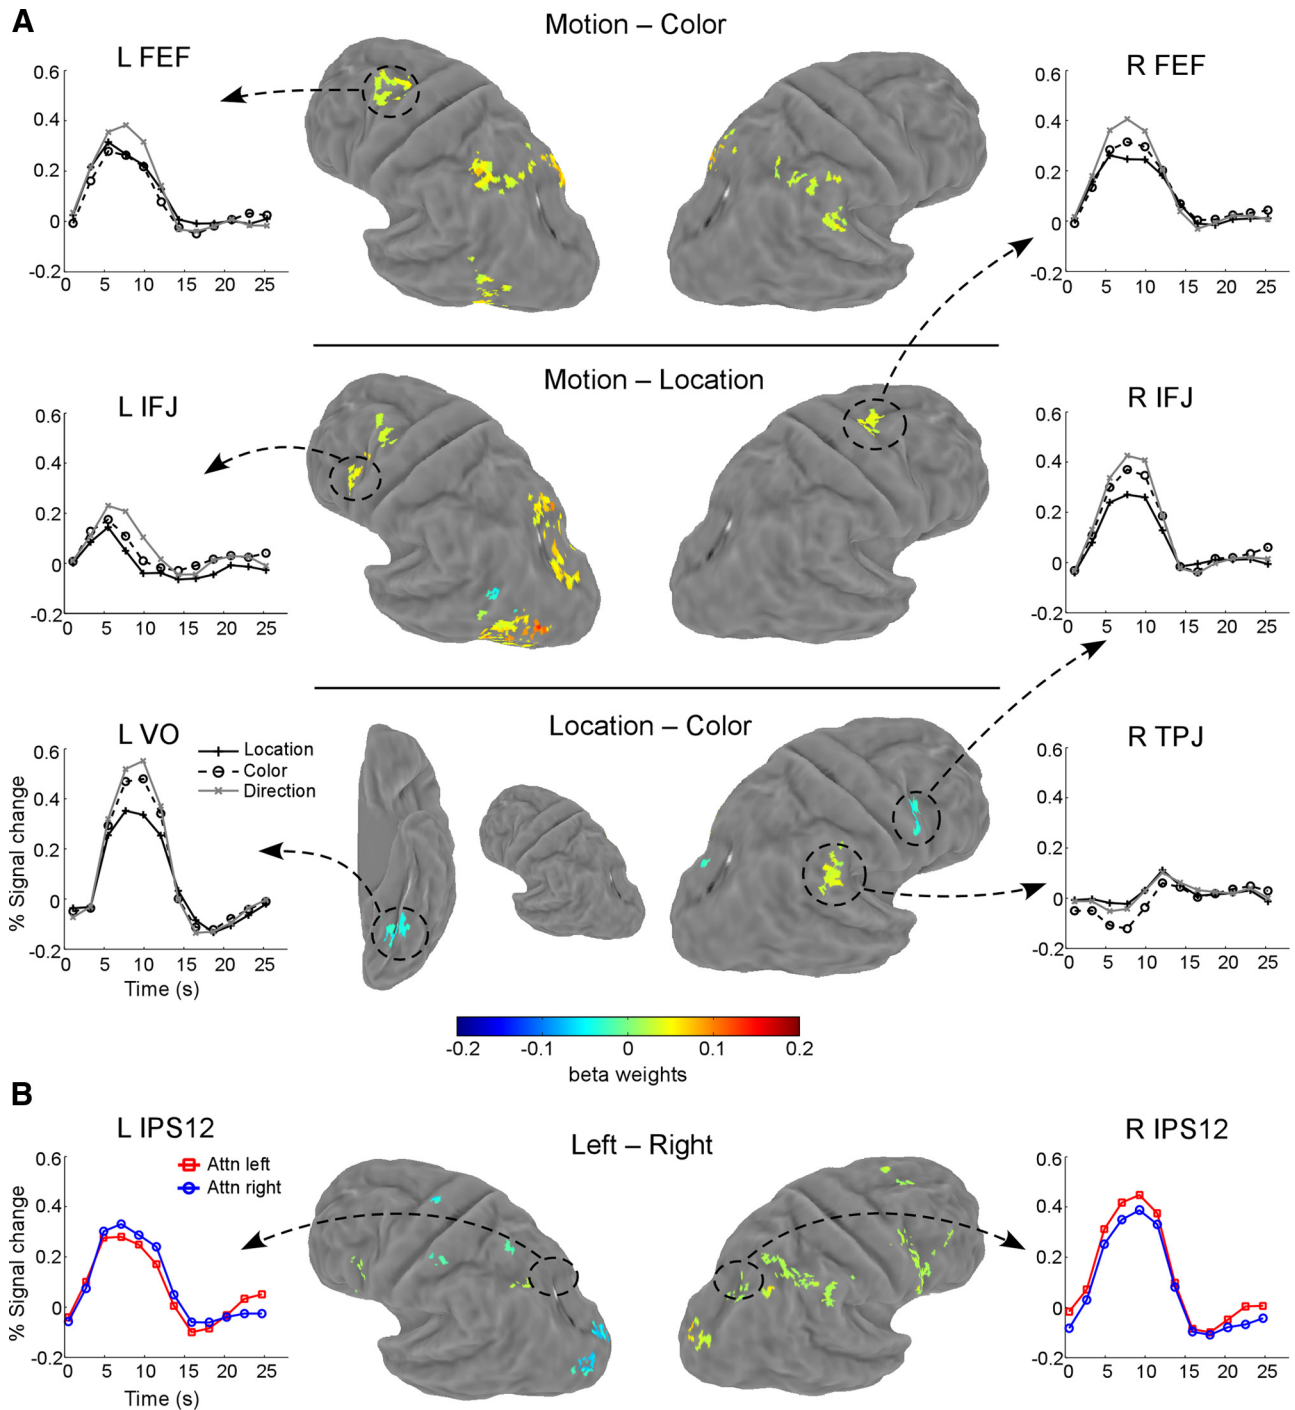


Figure 5. Maps of whole-brain contrasts. **A**, Comparison between attention types (location, color, and motion). Each row shows one contrast map on an atlas surface. Mean time courses in several select areas are also shown for illustration purpose only. **B**, Comparison between cues within each dimension. Shown here is a map for attention to left versus attention to right contrast. Time courses from the retinotopically defined IPS1/IPS2 region are shown to illustrate the effect of spatial attention. Contrasts between the two colors (red vs green) and two directions (up vs down) did not produce any significant activation. VO, Ventral occipital; TPJ: temporal parietal junction.

Whole-brain contrast analysis

The above univariate analyses were performed on predefined ROIs. We also performed whole-brain contrast analysis to search for voxels that were differentially activated by different types of attention. Contrast maps and time courses from select brain areas are shown in Figure 5A. Attention to motion tended to evoke a larger response than attention to color and location, and such differential activity was largely restricted to the dorsal frontoparietal cortex (Fig. 5A, first two rows). These results are consistent

with ROI-based analyses. In the Location–Color contrast (Fig. 5A, third row), we found three areas that showed differential activity. The right IFJ and an area in the left ventral occipital cortex (VO) showed higher fMRI response for the two feature attention conditions (motion and color) than the spatial attention condition (location). Another area in the right temporal parietal junction (TPJ) showed a larger initial decrease for attention to color than attention to motion and location. Overall results from the contrast analysis showed that attention to motion

evoked a larger response than attention to color and location, in a small set of cortical areas. We did not observe strong dissociations between attention conditions, e.g., a large response differential among conditions. In particular, we did not find any brain area that showed a higher fMRI response to spatial attention than feature-based attention.

We also performed contrast analyses between the cue pairs within each attention type (i.e., left vs right, red vs green, up vs down). With the same thresholding regime as the above analysis, we found a contralateral attentional effect in a small subset of frontoparietal areas for the left versus right contrast (Fig. 5B), but no significant activation for the red versus green and up versus down contrasts. To illustrate the effect of spatial attention, we also plotted the group-averaged time courses for retinotopically defined IPS1/IPS2 region, which showed contralateral attentional modulation.

Multivoxel similarity and clustering analyses

For each ROI in each participant, we correlated the multivoxel response vector for the six cue conditions to obtain a similarity matrix (see Materials and Methods). Mean similarity matrices across participants for select ROIs are shown in Figure 6. The entries on the main diagonal are the reliability values, a bootstrapped measure of the self-correlation of the response vectors, the median values of which are shown on top of each similarity matrix. In general, we obtained reliability values ~ 0.8 , suggesting a fairly stable pattern of multivoxel response vectors for a given cue condition. Off-diagonal entries in Figure 6 are pairwise correlation values (r) between response vectors for different cue conditions. In general, the reliability values were higher than pairwise correlations, as expected. Importantly, the off-diagonal values were not uniform, suggesting differential similarities across conditions.

To capture the structure in these similarity matrices, we performed clustering analyses, which were based on the distance between conditions, defined as $1 - r$ (see Materials and Methods). The results of the clustering analyses can be visualized by dendrograms, a tree structure with each cue condition as a node (Fig. 7). Nodes on the same branch were more similar (closer in distance) to each other than nodes on different branches. In the visual areas, there tended to be a cluster for spatial attention (attend to left and right) that was distinct from feature-based attention, but clustering within the latter was less organized. The auditory area showed some clustering of spatial attention and attention to motion, but the distances among all conditions were quite similar. In IPS12 and IPS34, color- and motion-based attention formed two distinct clusters, which tended to be different from spatial attention. The clearest clustering emerged in the two frontal areas, FEF and IFJ. In these areas, color, motion, and

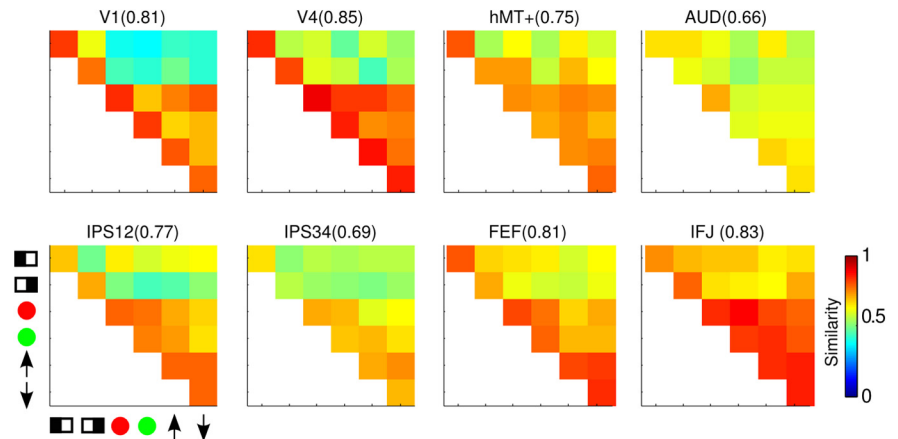


Figure 6. Mean similarity matrix across participants for select ROIs. Diagonal entries are the reliability of each condition and the median values are shown at the top of each part, in parenthesis after the name of the ROI. Off-diagonal entries are correlation values between each pair of cue condition, indicated by symbols on the horizontal and vertical axes (□, left; □, right; ●, red; ●, green; ↑, upward; ↓, downward).

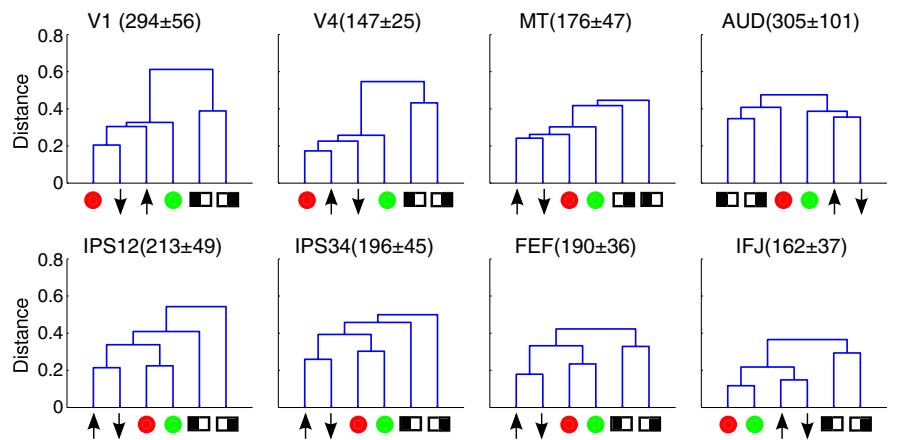


Figure 7. Dendrograms showing results from the clustering analysis. The vertical axis is the distance ($1 - r$) between different cue conditions. Symbols represented different cue conditions, in the same format as in Figure 6. The numbers in parenthesis are the average size of each ROI across participants, in number of voxels (mean \pm SD).

location all formed their own clusters. In addition, there was also a hierarchical structure among clusters, with the two feature dimensions forming a larger cluster, which was distinct from the cluster for spatial attention.

We first evaluated the reliability of the similarity and clustering results qualitatively, via permutation analyses (see Materials and Methods). When we randomly sampled half of the data, we obtained very similar results as using all data (Fig. 8A), but when we randomly sampled half of the data with shuffled trial labels, we did not find any apparent structure in the similarity matrix, which also led to drastically different dendrograms (Fig. 8B). These results indicate that the observed similarity and cluster structures were stable properties of the data.

To further quantitatively evaluate these results, we compared different groups of correlation coefficients via t tests. These comparisons were based on the hierarchical structure revealed by the clustering analyses. For each ROI, we partitioned the similarity matrix into five subsets, which is visualized in Figure 9A–C as follows: (1) correlations between the two spatial attention conditions (SS; Fig. 8A, magenta), (2) correlations between feature and spatial cues (fs; Fig. 8A, yellow), (3) correlations among all feature cues (ff; Fig. 8A, cyan), (4) correlations between two feature

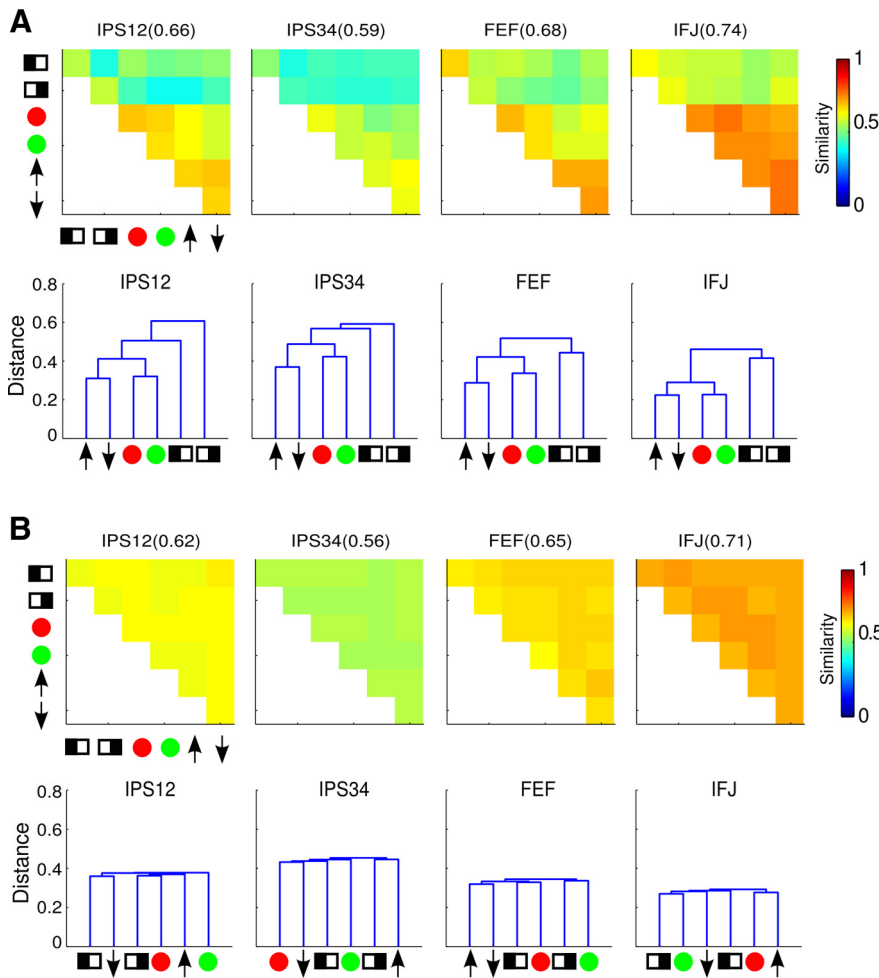


Figure 8. Results from resampling of trials. Only results for frontoparietal ROIs are shown; other ROIs exhibited similar results. **A**, Similarity matrices and dendrograms from randomly sampling half of the trials with original trial labels. **B**, Similarity matrices and dendrograms from randomly sampling half of the trials with shuffled trial labels.

cues within a dimension (wd; Fig. 8B, dark blue), (5) correlations between two feature cues across dimensions (bd; Fig. 8B, light blue). Note there were overlaps among partitions, such that the ff partition is composed of wd and bd partitions. We then averaged the correlation values within each partition, which we refer to as partition correlation, and compared these partition correlations to test for a significant hierarchical relationship.

We first tested for a separation between spatial attention and feature-based attention, by comparing partition correlations ss, fs, and ff (Fig. 9D). For all ROIs except AUD, we found a significantly higher correlation for the ff than the fs partition, indicating that the neural patterns were more similar to each other among attention to different features than between attention to features and locations. However, correlations for the ss and fs partition were similar, indicating neural patterns were no more similar between attention to the two locations (left and right) than between attention to locations and features. These results suggest a qualitative difference among conditions of feature-based attention and spatial attention. We then tested for a separation between attention to features within a dimension versus attention to features across dimensions, by comparing partition correlations wd and bd. These two partition correlations were equivalent in visual areas but wd correlation was significantly higher than bd correlation in frontoparietal cortex (IPS12, IPS34, FEF, IFJ). Thus, in these high-level cortical areas, attention to

features in the same dimension (e.g., red vs green) evoked more similar neural patterns than attention to features in different dimensions (e.g., red vs upward). These results suggest a further distinction between motion and color-based attention in these areas.

Discussion

In this study, we examined the relationship between attentional priority signals for different visual properties. Specifically, we assessed the pattern of neural signals in the dorsal attention network for spatial attention and two types of feature-based attention. Using multivoxel pattern similarity and clustering analysis, we found that neural signals in this network formed a hierarchical structure, such that spatial priority was distinct from feature priority, which in turn was further partitioned into dimension-specific priority signals (for motion and color). These results suggest that top-down attentional control is implemented by hierarchical neural signals that reflect the similarity structure of the task demands.

Our task design equated sensory stimulation as well as motor output among conditions, while only manipulating attentional instructions. Thus, any observed differences can be attributed to top-down attention, instead of sensorimotor aspects of the task. Our univariate analyses revealed that both spatial and feature-based attention activated very similar areas in the dorsal frontoparietal network (Wojciulik and Kanwisher, 1999; Giesbrecht et al., 2003; Slagter et al., 2007), supporting

the notion of a domain-general attentional control mechanism. Some of these studies contrasted spatial and feature-based attention and found that a subset of frontoparietal areas showed more activity for spatial cues than feature cues (Giesbrecht et al., 2003; Slagter et al., 2007). We also performed whole-brain contrast analyses but did not find preferential activation for spatial attention in dorsal frontoparietal areas (Fig. 5). Instead, we found attention to direction of motion tended to evoke the largest response in these areas. One important difference is that previous studies measured cue-related activity in a preparatory period without visual stimulus whereas we measured activity during active selection with a compound visual stimulus. It is possible that preparing for a potential target could rely on slightly different control mechanisms than actively maintaining priority for a visual target. In addition, there might be differences in participants’ ability to attend to the cued information without the visual stimulus. For example, it seems easier to attend to a spatial location than a color in an empty display, as space is always “present.” Previous studies were also limited in that they only compared spatial attention and color-based attention, whereas we also tested direction-based attention, in a design that equated selection demands and task difficulty. Thus, measured by average response amplitude, our data suggest a distinction between attention to static (color, locations) versus dynamic (motion)

properties, instead of between spatial and feature-based attention. Such a distinction is consistent with the finding that the dorsal frontal and parietal cortices contain many motion-sensitive regions (Culham et al., 2001; Orban et al., 2006). Areas with inherent motion sensitivity could exhibit a higher overall response during attention to motion even when the stimuli remained constant, as the motion signal presumably became more salient when attended.

This univariate analytical approach only provides a partial view of the data, however, as a contrast only compares the relative amplitude between two conditions. Thus, a statistically significant difference does not necessarily mean functional specificity. A complementary, and potentially more useful, approach in comparing different types of attentional priority is via multivariate analyses, which allows one to assess the *information content* carried by distributed patterns of neural activity (Kriegeskorte et al., 2006; Norman et al., 2006). Most current applications of fMRI multivariate analyses use the classifier approach, which assesses whether or not a particular task variable is represented by neural activities. Going beyond this type of categorical assessment, the multivoxel pattern similarity analysis combined with hierarchical clustering can further reveal the organizational structure of the underlying neural representations. Similar analytical techniques have been used in a previous study on sequential task control (Sigala et al., 2008). Our results demonstrate that neural signals in frontoparietal areas conform to a well organized hierarchical structure for attention to different properties. These results suggest that the frontoparietal areas contain a domain-general representation of attentional priority, where priority for all dimensions can be represented in the same neural populations, via different patterns. We should emphasize that our calculations of distance/similarity were not affected by differences in mean response amplitude (due to normalization of the response vector, see Materials and Methods), thus our similarity and clustering analyses provide new information that is not available from conventional univariate analyses.

Across all brain areas, we observed a more orderly clustering structure in the frontoparietal areas than occipital visual areas (Fig. 7). The visual areas exhibited two overall clusters between spatial and feature-based attention, whereas additional clustering along feature dimensions (motion vs color) was also observed in frontoparietal areas (Fig. 9). This gradient of more orderly clustering structure from visual to frontoparietal areas likely reflects the fact that activity in visual areas is driven more by bottom-up sensory input, which remained constant across conditions, whereas activity in frontoparietal areas is more tied to endogenous control, which varied with different attentional instructions. We also note that the AUD did not exhibit any clear clustering structure, which was expected given the acoustic features of the auditory cues did not have any built-in hierarchical structure. Thus results from AUD served as an internal control that also validated our analytical techniques.

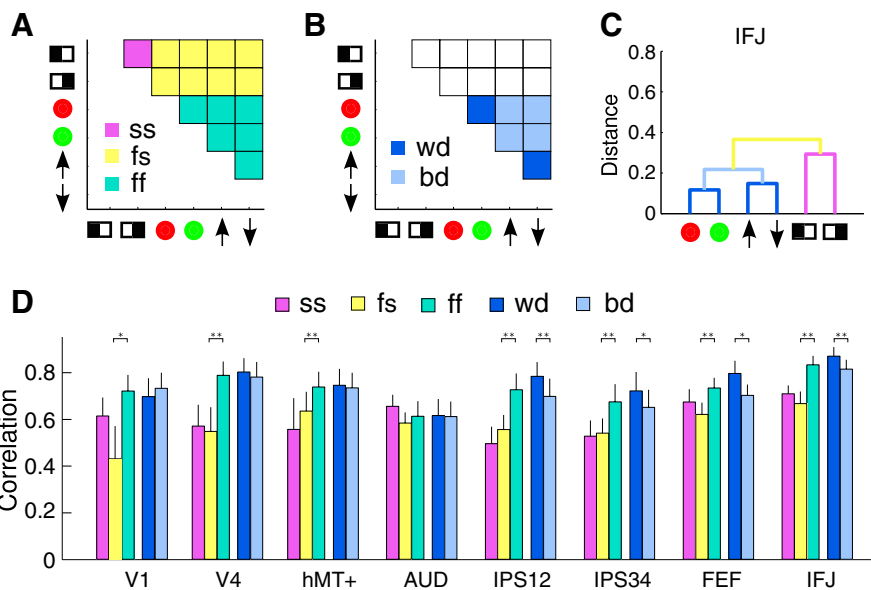


Figure 9. Quantitative evaluation of the hierarchical structure, with **A** and **B** showing the scheme of partitioning the similarity matrix for different comparisons. The similarity matrix is in the same format as in Figure 6, except that the main diagonal is removed. **A**, Shows the partition into spatial–spatial (SS), feature–feature (ff), and feature–spatial (fs) correlations. **B**, Shows the partition into within-dimension (wd) and between-dimension (bd) correlations. **C**, A color-coded dendrogram from area IFJ to illustrate the different pairs of conditions used in the partition scheme. The colors of the branches correspond to the color code in **A** and **B**. Note the ff (cyan) partition is not shown as it encompasses the wd (dark blue) and bd (light blue) partition. **D**, Average correlation values within each partition for select ROIs (partition correlations). The colors of the bars correspond to different partitions according to the color scheme shown in **A** and **B** (also shown in the legend). Error bars indicate ± 1 SEM across participants. Asterisks denote statistical significance between the two partitions for each comparison as evaluated via *t* tests ($*p < 0.05$; $**p < 0.005$).

Although pattern similarity between feature and spatial conditions was lower than that among different feature conditions (smaller fs than ff partition correlation; Fig. 9), suggesting a distinction between spatial and feature-based priority, the spatial–feature similarity was about the same as that between the two spatial conditions, left versus right (i.e., similar fs and ss partition correlation; Fig. 9). This was due to relatively low correlation levels between the two spatial conditions (Fig. 6, similarity matrix). This result likely reflects the contralateral nature of spatial priority, that is, attention to the left and right location primarily recruits brain areas in the right and left hemisphere, respectively. Thus there were less overlapped neural populations active in these two conditions, leading to a low correlation.

In addition to providing insights on neural representations of attentional priority, our results also have implications on the general question of the nature of internal representations. A prominent psychological theory, the so-called “second-order isomorphism,” proposes that internal representations should reflect the similarity relationship of physical stimuli (Shepard and Chipman, 1970). Computational analyses suggest that such a representation allows easy generalization to novel instances (Hinton et al., 1986; Edelman, 1998). Neuroimaging studies of object perception have found support of this idea in neural representations of object shape in the ventral visual stream (Kiani et al., 2007; Kriegeskorte et al., 2008; Weber et al., 2009). Our results can be interpreted as an extension of this concept to the domain of attentional control in the dorsal stream. That is, the underlying neural representation for attentional priority conforms to the similarity structure of the selection demands. Thus, the fact that priority signals are more similar between different colors than between colors and locations can afford the system to encode

priority for arbitrary colors and locations (i.e., generalization). Such an organization could support the flexibility in attentional selection.

Our results demonstrate that different types of attentional priority can be represented by distinct multivoxel patterns in the same cortical area. This observation suggests a flexible underlying neuronal representation where the same neurons can represent attended features in multiple dimensions and can dynamically adjust their feature tuning depending on the task. Such a scenario is reminiscent of studies that showed the same prefrontal neurons can represent different stimulus attributes in working memory tasks (Rao et al., 1997). In general, theoretical analyses have suggested that prefrontal cortex contains highly flexible neural codes for task variables (Duncan, 2001; Miller and Cohen, 2001), which is supported by recent studies using tasks more complex than the selection task in our study (Li et al., 2007; Woolgar et al., 2011a,b). In a similar vein, attentional priority could be represented by frontoparietal neurons tuned to multiple dimensions, i.e., signal multiplexing at the single neuron level. This conjecture can be tested in future studies using single-unit recording techniques.

In conclusion, our data revealed distinct neural patterns underlying different types of attentional selection. These neural patterns formed a systematic, hierarchical structure in the dorsal frontoparietal network that reflects the similarity of task demands. Such an organization likely affords the system to represent attentional priority in a flexible manner. Multivariate analyses provided unique insight on the neural organization of priority signals, which is not obtainable from conventional univariate analyses. Our data further point at signal multiplexing as a possible neuronal mechanism for the representation of attentional priority.

References

- Baldassi S, Verghese P (2005) Attention to locations and features: different top-down modulation of detector weights. *J Vis* 5(6):556–570. [Medline](#)
- Bisley JW, Goldberg ME (2010) Attention, intention, and priority in the parietal lobe. *Annu Rev Neurosci* 33:1–21. [CrossRef Medline](#)
- Botvinick M, Plaut DC (2004) Doing without schema hierarchies: a recurrent connectionist approach to normal and impaired routine sequential action. *Psychol Rev* 111:395–429. [CrossRef Medline](#)
- Brewer AA, Liu J, Wade AR, Wandell BA (2005) Visual field maps and stimulus selectivity in human ventral occipital cortex. *Nat Neurosci* 8:1102–1109. [CrossRef Medline](#)
- Carrasco M (2006) Covert attention increases contrast sensitivity: psychophysical, neurophysiological, and neuroimaging studies. In: *Visual perception* (Martinez-Conde S, Macknik SL, Martinez LM, Alonso JM, Tse PU, eds). Amsterdam: Elsevier.
- Carrasco M (2011) Visual attention: the past 25 years. *Vision Res* 51:1484–1525. [CrossRef Medline](#)
- Cooper RP, Shallice T (2006) Hierarchical schemas and goals in the control of sequential behavior. *Psychol Rev* 113:887–916; discussion 971–831. [CrossRef Medline](#)
- Corbetta M, Shulman GL (2002) Control of goal-directed and stimulus-driven attention in the brain. *Nat Rev Neurosci* 3:201–215. [Medline](#)
- Culham J, He S, Dukelow S, Verstraten FA (2001) Visual motion and the human brain: what has neuroimaging told us? *Acta Psychol* 107:69–94. [CrossRef Medline](#)
- Desimone R, Duncan J (1995) Neural mechanisms of selective visual attention. *Annu Rev Neurosci* 18:193–222. [CrossRef Medline](#)
- DeYoe EA, Carman GJ, Bandettini P, Glickman S, Wieser J, Cox R, Miller D, Neitz J (1996) Mapping striate and extrastriate visual areas in human cerebral cortex. *Proc Natl Acad Sci U S A* 93:2382–2386. [CrossRef Medline](#)
- Duncan J (2001) An adaptive coding model of neural function in prefrontal cortex. *Nat Rev Neurosci* 2:820–829. [CrossRef Medline](#)
- Edelman S (1998) Representation is representation of similarities. *Behav Brain Sci* 21:449–467; discussion 467–498. [Medline](#)
- Engel SA, Glover GH, Wandell BA (1997) Retinotopic organization in human visual cortex and the spatial precision of functional MRI. *Cereb Cortex* 7:181–192. [CrossRef Medline](#)
- Gardner JL, Sun P, Waggoner RA, Ueno K, Tanaka K, Cheng K (2005) Contrast adaptation and representation in human early visual cortex. *Neuron* 47:607–620. [CrossRef Medline](#)
- Giesbrecht B, Woldorff MG, Song AW, Mangun GR (2003) Neural mechanisms of top-down control during spatial and feature attention. *Neuroimage* 19:496–512. [CrossRef Medline](#)
- Guo F, Preston TJ, Das K, Giesbrecht B, Eckstein MP (2012) Feature-independent neural coding of target detection during search of natural scenes. *J Neurosci* 32:9499–9510. [CrossRef Medline](#)
- Hinton GE, McClelland JL, Rumelhart DE (1986) Distributed representations. In: *Parallel distributed processing: explorations in the microstructure of cognition. Volume I: foundations.* (Rumelhart DE, McClelland JL, PDP Research Group eds), pp 77–109. Cambridge, MA: MIT.
- Hou Y, Liu T (2012) Neural correlates of object-based attentional selection in human cortex. *Neuropsychologia* 50:2916–2925. [CrossRef Medline](#)
- Kaiser PK (1991) Flicker as a function of wavelength and heterochromatic flicker photometry. In: *Limits of vision* (Kulikowski JJ, Walsh V, Murray IJ, eds), pp 171–190. Basingstoke: Macmillan.
- Kastner S, Ungerleider LG (2000) Mechanisms of visual attention in the human cortex. *Annu Rev Neurosci* 23:315–341. [CrossRef Medline](#)
- Kiani R, Esteky H, Mirpour K, Tanaka K (2007) Object category structure in response patterns of neuronal population in monkey inferior temporal cortex. *J Neurophysiol* 97:4296–4309. [CrossRef Medline](#)
- Kriegeskorte N, Goebel R, Bandettini P (2006) Information-based functional brain mapping. *Proc Natl Acad Sci U S A* 103:3863–3868. [CrossRef Medline](#)
- Kriegeskorte N, Mur M, Ruff DA, Kiani R, Bodurka J, Esteky H, Tanaka K, Bandettini PA (2008) Matching categorical object representations in inferior temporal cortex of man and monkey. *Neuron* 60:1126–1141. [CrossRef Medline](#)
- Li S, Ostwald D, Giese M, Kourtzi Z (2007) Flexible coding for categorical decisions in the human brain. *J Neurosci* 27:12321–12330. [CrossRef Medline](#)
- 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. [CrossRef Medline](#)
- Liu T, Stevens ST, Carrasco M (2007) Comparing the time course and efficacy of spatial and feature-based attention. *Vision Res* 47:108–113. [CrossRef Medline](#)
- Liu T, Hospadaruk L, Zhu DC, Gardner JL (2011) Feature-specific attentional priority signals in human cortex. *J Neurosci* 31:4484–4495. [CrossRef Medline](#)
- Martinez-Trujillo JC, Treue S (2004) Feature-based attention increases the selectivity of population responses in primate visual cortex. *Curr Biol* 14:744–751. [CrossRef Medline](#)
- Miller EK, Cohen JD (2001) An integrative theory of prefrontal cortex function. *Annu Rev Neurosci* 24:167–202. [CrossRef Medline](#)
- Moore T (2006) The neurobiology of visual attention: finding sources. *Curr Opin Neurobiol* 16:159–165. [CrossRef Medline](#)
- Nichols TE, Holmes AP (2002) Nonparametric permutation tests for functional neuroimaging: a primer with examples. *Hum Brain Mapp* 15:1–25. [CrossRef Medline](#)
- Norman KA, Polyn SM, Detre GJ, Haxby JV (2006) Beyond mind-reading: multivoxel pattern analysis of fMRI data. *Trends Cogn Sci* 10:424–430. [CrossRef Medline](#)
- Nunnally JC (1978) *Psychometric theory*. New York: McGraw-Hill.
- Orban GA, Claeys K, Nelissen K, Smans R, Sunaert S, Todd JT, Wardak C, Durand JB, Vanduffel W (2006) Mapping the parietal cortex of human and nonhuman primates. *Neuropsychologia* 44:2647–2667. [CrossRef Medline](#)
- Paus T (1996) Location and function of the human frontal eye-field: a selective review. *Neuropsychologia* 34:475–483. [CrossRef Medline](#)
- Rao SC, Rainer G, Miller EK (1997) Integration of what and where in the primate prefrontal cortex. *Science* 276:821–824. [CrossRef Medline](#)
- Reynolds JH, Chelazzi L (2004) Attentional modulation of visual processing. *Annu Rev Neurosci* 27:611–647. [CrossRef Medline](#)
- Saenz M, Buracas GT, Boynton GM (2002) Global effects of feature-based attention in human visual cortex. *Nat Neurosci* 5:631–632. [CrossRef Medline](#)

- Serences JT, Yantis S (2006) Selective visual attention and perceptual coherence. *Trends Cogn Sci* 10:38–45. [CrossRef Medline](#)
- Sereno MI, Dale AM, Reppas JB, Kwong KK, Belliveau JW, Brady TJ, Rosen BR, Tootell RB (1995) Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Science* 268:889–893. [CrossRef Medline](#)
- Shepard RN, Chipman S (1970) Second-order isomorphism of internal representations: shapes of states. *Cognitive Psychol* 1:1–17. [CrossRef](#)
- Sigala N, Kusunoki M, Nimmo-Smith I, Gaffan D, Duncan J (2008) Hierarchical coding for sequential task events in the monkey prefrontal cortex. *Proc Natl Acad Sci U S A* 105:11969–11974. [CrossRef Medline](#)
- Slagter HA, Giesbrecht B, Kok A, Weissman DH, Kenemans JL, Woldorff MG, Mangun GR (2007) fMRI evidence for both generalized and specialized components of attentional control. *Brain Res* 1177:90–102. [CrossRef Medline](#)
- Swisher JD, Halko MA, Merabet LB, McMains SA, Somers DC (2007) Visual topography of human intraparietal sulcus. *J Neurosci* 27:5326–5337. [CrossRef Medline](#)
- Tan PN, Steinbach M, Kumar V (2005) *Introduction to data mining*. Boston: Addison-Wesley.
- Todd MT, Nystrom LE, Cohen JD (2013) Confounds in multivariate pattern analysis: theory and rule representation case study. *Neuroimage* 77:157–165. [CrossRef Medline](#)
- Van Essen DC (2005) A population-average, landmark- and surface-based (PALS) atlas of human cerebral cortex. *Neuroimage* 28:635–662. [CrossRef Medline](#)
- Wandell BA, Dumoulin SO, Brewer AA (2007) Visual field maps in human cortex. *Neuron* 56:366–383. [CrossRef Medline](#)
- Watson JD, Myers R, Frackowiak RS, Hajnal JV, Woods RP, Mazziotta JC, Shipp S, Zeki S (1993) Area V5 of the human brain: evidence from a combined study using positron emission tomography and magnetic resonance imaging. *Cereb Cortex* 3:79–94. [CrossRef Medline](#)
- Weber M, Thompson-Schill SL, Osherson D, Haxby J, Parsons L (2009) Predicting judged similarity of natural categories from their neural representations. *Neuropsychologia* 47:859–868. [CrossRef Medline](#)
- Wojciulik E, Kanwisher N (1999) The generality of parietal involvement in visual attention. *Neuron* 23:747–764. [CrossRef Medline](#)
- Wolfe JM (1994) Guided Search 2.0: a revised model of visual search. *Psychon Bull Rev* 1:202–238. [CrossRef](#)
- Woolgar A, Hampshire A, Thompson R, Duncan J (2011a) Adaptive coding of task-relevant information in human frontoparietal cortex. *J Neurosci* 31:14592–14599. [CrossRef Medline](#)
- Woolgar A, Thompson R, Bor D, Duncan J (2011b) Multi-voxel coding of stimuli, rules, and responses in human frontoparietal cortex. *Neuroimage* 56:744–752. [CrossRef Medline](#)