

Cognition and Behavior

Neural Determinants of Task Performance during Feature-Based Attention in Human Cortex

Michael Jigo,^{1,2} Mengyuan Gong,¹ and Taosheng Liu¹DOI:<http://dx.doi.org/10.1523/ENEURO.0375-17.2018>¹Department of Psychology, Michigan State University, East Lansing, MI 48824 and ²Center for Neural Science, New York University, New York, NY 10003

Abstract

Studies of feature-based attention have associated activity in a dorsal frontoparietal network with putative attentional priority signals. Yet, how this neural activity mediates attentional selection and whether it guides behavior are fundamental questions that require investigation. We reasoned that endogenous fluctuations in the quality of attentional priority should influence task performance. Human subjects detected a speed increment while viewing clockwise (CW) or counterclockwise (CCW) motion (baseline task) or while attending to either direction amid distracters (attention task). In an fMRI experiment, direction-specific neural pattern similarity between the baseline task and the attention task revealed a higher level of similarity for correct than incorrect trials in frontoparietal regions. Using transcranial magnetic stimulation (TMS), we disrupted posterior parietal cortex (PPC) and found a selective deficit in the attention task, but not in the baseline task, demonstrating the necessity of this cortical area during feature-based attention. These results reveal that frontoparietal areas maintain attentional priority that facilitates successful behavioral selection.

Key words: covert attention; fMRI; frontoparietal cortex; multivariate analysis; TMS

Significance Statement

To cope with the computational limits of visual processing, the brain selectively prioritizes a subset of visual input. The selection of visual features, such as color and motion, has been associated with activity in a frontoparietal cortical network. Yet, the role this activity plays in mediating selection and influencing behavior is not clear. Using fMRI, we show that neural activity patterns in several frontoparietal areas correlated with task performance. Furthermore, neurodisruption of the posterior parietal cortex (PPC) using transcranial magnetic stimulation (TMS) selectively impaired feature selection. These results provide the first evidence that the neural representation of prioritized features in frontoparietal areas play a causal role in selecting visual features.

Introduction

Visual attention allows us to select relevant information from the visual scene for prioritized processing. Given the

profound effect of attention on perception (Simons and Rensink, 2005), it is important to understand how attention is controlled both at the behavioral and neural level (Corbetta and Shulman, 2002).

Received November 3, 2017; accepted February 8, 2018; First published February 19, 2018.

The authors declare no competing financial interests.

Author contributions: T.L. and M.J. designed research; M.J., T.L., and M.G. performed research; M.J., T.L., and M.G. analyzed data; M.J., M.G., and T.L. wrote the paper.

This work was supported by the National Institutes of Health Grant R01EY022727.

Acknowledgements: We thank Dr. David Zhu and Ms. Scarlett Doyle for their assistance in collecting the neuroimaging data; Rebecca Francis, Samantha Blair, and Chris Kmiec for their assistance in collecting the magnetic stimulation data; and Dr. Samuel Ling for his comments on an earlier draft of this manuscript.

Correspondence should be addressed to Taosheng Liu, Department of Psychology, Michigan State University, 316 Physics Road, East Lansing, MI 48864, E-mail: tsliu@msu.edu.

Theories of attention have postulated that attentional priority, which represents the behavioral importance of visual objects, guides selection during perceptual processing (Wolfe, 1994; Desimone and Duncan, 1995; Deco and Rolls, 2004). On the neural level, attentional priority signals have been linked to activity in dorsal parietal and frontal areas. In particular, priority for spatial locations is thought to be represented by spatially selective neural responses in these areas (Silver and Kastner, 2009; Bisley and Goldberg, 2010). This idea has been strongly supported by neurophysiological and neuroimaging studies (for review, see Buschman and Kastner, 2015). In addition to spatial locations, attention can also be directed to nonspatial properties such as visual features (Yantis, 2000), and recent work has begun to characterize priority signals for feature-based attention (Serences and Boynton, 2007; Liu et al., 2011; Liu and Hou, 2013; Ester et al., 2016). Specifically, these studies have found that patterns of neural activity in visual and dorsal frontoparietal areas, in particular, areas along the intraparietal sulcus and the precentral sulcus, can be used to decode the attended feature value (e.g., red vs green color, leftward vs rightward motion). Although these studies demonstrate that neural activity patterns vary with the attended feature, they do not reveal how this activity mediates attentional selection and whether it is causally involved in feature selection and task performance.

Here, we investigate the neural-behavioral relation in regions that have been implicated in the selection of visual features. We reasoned that if neural activity encodes attentional priority for visual features, then such activity should be related to behavioral performance in a task requiring feature-based selection. In an fMRI experiment, we instructed human subjects to attend to one of two superimposed motion directions and perform a speed detection at psychophysical threshold (attention task). If neural activity in this task encodes priority for features, the quality of such signals should be better for correct than incorrect trials. To provide a benchmark to evaluate the quality of the priority signal, we measured neural activity in a separate task where subjects viewed a single motion direction (baseline task). We hypothesized that if attentional selection in the attention task is successful, then neural activity will resemble that in the baseline task, inasmuch as attention reduces the influence from the distracter. Conversely, neural activity for unsuccessful selection in the attention task would share less resemblance to that in the baseline task. Thus, we predict that the neural pattern similarity between the attention and baseline task would vary with task performance.

To assess the causal role of neural activity on feature-based selection, we used transcranial magnetic stimulation (TMS) to disrupt identified neural signals for attentional priority. We further hypothesized that if these

signals represent priority for features, then neurodisruption would impair performance in the attention task, which requires feature-based selection, whereas it would impair performance less (or not at all) in the baseline task, which does not require feature-based selection. Conversely, we hypothesized that if these signals are related to general motion processing, then neurodisruption would produce equivalent impairments in both the attention and baseline task. To test these predictions, we unilaterally targeted representative brain areas that were identified as potential loci for maintaining attentional priority or stimulus processing.

Materials and Methods

fMRI experiment

Subjects

Twelve subjects participated in the imaging experiment (five males, 20–29 years old). All subjects were neurologically intact, had normal or corrected-to-normal vision, and were recruited from the Michigan State University community (undergraduate and graduate students). They gave written informed consent under the study protocol approved by the Institutional Review Board at Michigan State University and were remunerated at a rate of \$20 per hour.

Visual display and stimuli

Visual stimuli were generated using MGL (<http://gru.stanford.edu/doku.php/mgl/overview>), a set of OpenGL libraries running in Matlab (Mathworks). In the psychophysics laboratory, stimuli were presented on a 19" CRT monitor (resolution: 1024 × 768, 60-Hz refresh rate) and subjects had their heads stabilized by a chin rest that was positioned 85 cm away from the monitor. During MR scanning sessions, a DLP projector (Psychology Software Tools) projected the stimuli onto a rear-projection screen located in the scanner bore. Subjects viewed the screen via an angled mirror attached to the head coil at a viewing distance of 60 cm. The projector had a resolution of 1024 × 768 and was updated at 60 Hz.

The stimuli were composed of one or two dot fields that rotated in the clockwise (CW) or counterclockwise (CCW) direction with 60% motion coherence. Each dot field was contained within an annulus (eccentricity from 2.5° to 8°) that was centered on a central cross (size: 0.5°) and displayed on a black background. Each dot within a field (dot color: gray; dot size: 0.1°; density: 1.1 dots/deg²) had a lifetime of six frames (0.1 s) to deter subjects from tracking individual dots. During training and imaging sessions, the central cross was either yellow or cyan to help subjects remember the current response mapping (see *Tasks*).

Tasks

Attention task

At trial onset, subjects were cued to attend to CW or CCW motion by a rightward or leftward-pointing arrow cue, respectively (Fig. 1A). The cue appeared 0.77° above the central cross and persisted on the display for 0.3 s. The cue was then replaced with spatially overlapping CW and CCW dot fields that were displayed for 4.1 s. Each

DOI:<http://dx.doi.org/10.1523/ENEURO.0375-17.2018>

Copyright © 2018 Jigo et al.

This is an open-access article distributed under the terms of the [Creative Commons Attribution 4.0 International license](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution and reproduction in any medium provided that the original work is properly attributed.

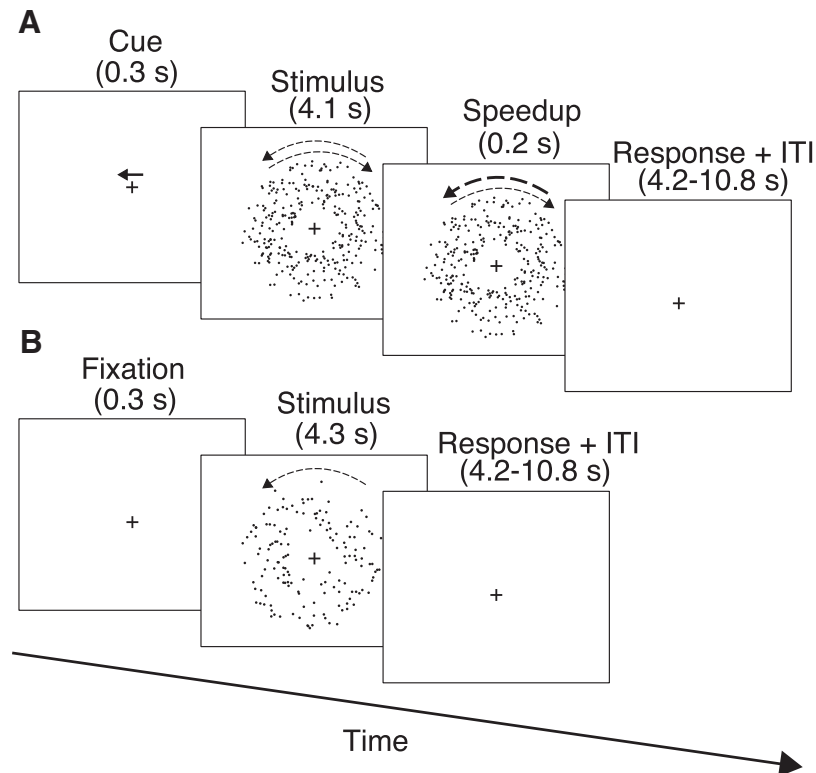


Figure 1. Schematic of the attention and baseline tasks. **A**, Sequence of a valid trial in the attention task. Size of curved arrows illustrates the speed of rotation. The fixation cross is either yellow or cyan (color not shown). **B**, Sequence of a target-absent trial in the baseline task. For ease of illustration, frames depict black stimuli on a white background (colors are reversed in the actual experiment).

dot field rotated around the center of the annulus at a speed of $45^\circ/\text{s}$ before a brief (0.2 s) speed increment occurred in either direction. Subjects were instructed to maintain fixation on the central cross and report whether the speedup occurred in the cued direction. Thus, task performance was contingent on the attended direction. If CW (CCW) was cued when a CW (CCW)-speedup occurred, and the subject reported the speedup, the trial was classified as a hit. Alternatively, if the subject failed to report the speedup, the trial was classified as a miss. On 80% of trials, the speedup occurred in the cued direction (valid trials) and its magnitude was adjusted via best Parameter Estimation by Sequential Testing (PEST), an adaptive staircase procedure as implemented in the Palamedes toolbox (Prins and Kingdon, 2009), to maintain a hit rate (performance) of 65%. The best PEST procedure computed the maximum-likelihood estimate of an observer's speed increment threshold on each trial, on the basis of all previous responses. Performance was assumed to have the form of a Weibull function, with the slope fixed at 2, and the lapse rate and guess rate fixed at 1%. On invalid trials (20% of trials), a speedup occurred in the uncued direction (invalid trials) using the magnitude of the preceding valid trial (i.e., the speedup on invalid trials was not controlled via staircase).

For the attention task, we aimed to maximize the number of hit and miss trials to conduct our multivariate analyses. To accomplish this, we used a 4:1 validity ratio and titrated performance to an intermediate hit rate (65%).

Although a lower hit rate, such as 50%, would mathematically maximize the number of hits and misses, such a low performance level could discourage subjects from using the cue.

An intertrial interval (ITI) followed the speedup, the ITI was 4.2 s on 40% of trials, 6.4 s on 30% of trials, 8.6 s on 20% of trials, and 10.8 s on 10% of trials. During this interval, subjects reported the presence or absence of a speedup in the cued direction via a “present” or “absent” response that was mapped onto a particular finger. To differentiate the observed neural response for target detection from that of a motor plan, we trained subjects to use two inverse response mappings that were indicated by the color of the central cross. When the cross was cyan, subjects made present and absent responses with their index and middle finger, respectively. When the cross was yellow, the mapping was reversed. The response mappings alternated across runs and were counterbalanced within each subject.

Baseline task The baseline task (Fig. 1B) was identical to the attention task, with the following exceptions: (1) an arrow cue was not presented, hence, subjects viewed the central cross for 0.3 s at the beginning of each trial; (2) only one dot field was displayed; (3) a speedup occurred on 70% of trials (target-present trials) and on the other 30% of trials (target-absent trials), no speedup occurred; and (4) the hit rate was maintained at 50% via best PEST. These behavioral manipulations were used to maximize the number of correct rejections (when subjects reported

absent on target-absent trials), as these served as the baseline neural patterns in our multivariate analyses. In particular, the low hit rate was chosen to dissuade subjects from making too many false positive reports, which would reduce the correct rejection rate.

Procedure

Training session Before the imaging session, subjects completed at least two runs (50 trials/run) of each task in the psychophysics laboratory for practice. To ensure that subjects maintained fixation on the central cross, their left eye was recorded at a sampling rate of 1000 Hz using an Eyelink 1000 system (SR Research); the data were analyzed offline using custom Matlab scripts. Importantly, subjects were only allowed to proceed with the imaging session if their eye position was always within 2° of the central cross.

Because speedup events only occurred at the end of a trial, this could allow for a strategy where attention was only directed to the stimulus later in the trial. Thus, we also emphasized the importance of maintaining attention throughout each trial and all subjects reported compliance. We note that partial and/or inconsistent attentional engagement would likely reduce the attentional signal and increase noise in neural data. Hence, our observed effects might be even stronger if attention had been engaged more consistently throughout the stimulus duration.

Imaging session Before functional images were collected, subjects completed, at most, 100 trials of each task while they lay in the scanner; these trials were used to calibrate the staircase to maintain the expected hit rate. Then, subjects completed 12 fMRI runs (30 trials/run), with six runs for each task in an alternating sequence. Each run began with an 8.8-s fixation period and lasted 338.8 s; the images collected during the fixation period were discarded to avoid magnetic saturation effects. For the attention task, cue direction (CW vs CCW) and validity (valid vs invalid) were randomly interleaved within each run; whereas, during the baseline task, motion direction (CW vs CCW) and trial type (target-present vs target-absent) were randomly interleaved within each run. Eye position was not monitored during this session.

MRI data acquisition

Imaging was performed on a GE Healthcare 3 T Sigma HDx MRI scanner, equipped with an eight-channel head coil, in the Department of Radiology at Michigan State University. For each subject, high-resolution anatomic images were acquired using a T1-weighted magnetization-prepared rapid-acquisition gradient echo sequence (field of view, 256 × 256 mm; 180 sagittal slices; 1-mm isotropic voxels). Functional images were acquired using a T2*-weighted echo planar imaging sequence (repetition time, 2.2 s; echo time, 30 ms; flip angle, 78°; matrix size, 64 × 64; in-plane resolution, 3 × 3 mm; slice thickness, 4 mm, interleaved, no gap). Thirty axial slices covering the whole brain were collected. In each scanning session, we also acquired a 2D T1-weighted anatomic image that had the same slice prescription as the functional scans but with higher in-plane resolution (0.75 × 0.75 × 4 mm). This

image was used to align the functional data to the high-resolution anatomic images for each subject.

Retinotopic mapping

In an independent scanning session, we mapped each subject's occipital visual cortex and several parietal areas that contain topographic maps. For the occipital cortex, we used rotating wedges and expanding/contracting rings (eccentricity from 0.5° to 8.25°) to map the polar angle and radial component, respectively (Serenó et al., 1995; DeYoe et al., 1996; Engel et al., 1997). Four runs of the wedge stimuli and two runs of the ring stimuli were collected and averaged. A Fourier analysis was then applied to the averaged time course to derive the amplitude and phase of the response, the latter forming the polar angle map of the responses. Borders between areas were defined as phase reversals in the polar angle map of the visual field. The map was visualized on computationally flattened representations of the cortical surface, which were generated from a high-resolution anatomic image using FreeSurfer (<http://surfer.nmr.mgh.harvard.edu>) and custom Matlab code.

Parietal areas were mapped with a memory delayed saccade task that was modeled after previous studies on parietal topography (Serenó et al., 2001; Schluppeck et al., 2005; Konen and Kastner, 2008). Subjects fixated on a central point while a peripheral (~10° radius) target dot was flashed for 500 ms. The flashed peripheral stimulus was quickly replaced by a ring of 100 distractor dots randomly positioned within a ring with a radius of 8.5–10.5°. The distractors remained on screen for 3 s, after which subjects made a saccade to the remembered position of the peripheral target and then immediately made a saccade back to the central fixation point. The position of the peripheral saccade target shifted around the periphery from trial to trial in either a CW or CCW direction, so that after eight trials the target completed one full cycle. A trial lasted 5 s and six cycles were completed in a single run. Two to four runs of the memory delayed saccade task were collected and averaged, then borders between parietal areas were defined as phase reversals in the polar angle map.

Finally, we presented moving versus stationary dots (eccentricity from 0.5° to 10°) in alternating blocks and localized the human motion-sensitive area as an area near the junction of the occipital and temporal cortex that responded more to moving than stationary dots (Watson et al., 1993). This area likely contained both MT and MST, so we refer to it as MT+.

Overall, we identified the following areas in each subject: V1, V2, V3, V3A/B, V4, MT+, V7/IPS0, IPS1, and IPS2. We did not observe a consistent boundary between V3A and V3B; hence, we defined an area that contained both and labeled it V3A/B. We adopted the definition of V4 as a hemifield representation anterior to V3v (Wandell et al., 2007). The V7/IPS0 nomenclature was adopted because its anatomic location is within the IPS in some hemispheres and shares a foveal representation with IPS1 (Swisher et al., 2007). We could not reliably observe borders for more anterior IPS regions such as IPS3 and IPS4 in all subjects.

fMRI data analysis

Preprocessing Functional MRI data were analyzed using mrTools (<http://gru.stanford.edu/doku.php/mr-Tools/overview>) running in Matlab and custom code in Matlab. Preprocessing of functional data included head motion correction, linear detrending, and temporal high pass filtering at 0.01 Hz. The 2D T1-weighted image was used to compute the alignment between the functional images and the high-resolution T1-weighted image, using an automated robust image registration algorithm (Nestares and Heeger, 2000). Functional data were converted to percentage signal change by dividing the time course of each voxel by its mean signal over a run. Then, data for all runs of a task were concatenated, resulting in two time series. All region of interest (ROI) analyses were performed in individual subject's native anatomic space.

Univariate analysis: deconvolution For the attention task, we fit each voxel's time series with a general linear model containing five sets of regressors: four corresponding to the two directions during valid trials (CW vs CCW) crossed by response accuracy (hits vs misses) and the fifth corresponding to invalid trials. We note that due to the low proportion of invalid trials, false alarms were rather scarce (5 ± 5 false alarms across subjects), precluding a further separation into CW and CCW trials. In the analyses described below (see *Multivariate analysis*), we focus only on valid trials in the attention task; hence, hits and miss trials are referred to as correct and incorrect trials, respectively. For the baseline task, the general linear model contained seven sets of regressors: six corresponding to the two motion directions (CW vs CCW) crossed by detection type (hit vs miss vs correct rejection) and the seventh corresponding to false alarms. Each regressor modeled the fMRI response in a 17.6-s window after trial onset with a set of finite impulse responses. The design matrix was pseudo-inversed and multiplied by the time series to obtain an estimate of the hemodynamic response for each condition (deconvolution).

To obtain precise estimates of BOLD response amplitude for each subject, we averaged their deconvolved response across ROIs to obtain an overall response profile. These response profiles revealed a variable time-to-peak across subjects that ranged from 4.4 to 8.8 s after trial onset (time points 3–5). Voxel-wise estimates of response amplitude for each condition were then computed as the average deconvolved response from the time point immediately preceding the subject's peak to the time point following the peak.

From the deconvolution model of each task, we obtained a goodness of fit measure (r^2 value) for each voxel, which was the amount of variance in the fMRI time series accounted for by the model (Gardner et al., 2005). The r^2 value indicated how much a voxel's time series was driven by the task. For each subject, we used their r^2 map of the attention task to define two frontal ROIs. These ROIs were centered on voxels with maximal r^2 values that formed separate clusters along the precentral sulcus: one near the superior frontal sulcus and another near the inferior frontal sulcus. These clusters are not separate on the group map (Fig. 3A) but were distinct in individual

subjects. At each location, we defined a ROI that extended to enclose the nearby sulcal junction while avoiding the precentral gyrus (motor cortex). The dorsal ROI coincided with the putative human frontal eye field (FEF) and we referred to the ventral ROI as the inferior frontal junction (IFJ); both areas were defined bilaterally.

Multivariate analysis: correlation To measure attentional priority, we calculated the correlation between BOLD response patterns when subjects attended to a feature (attention task) and when they viewed the feature in isolation (baseline task; Fig. 4). The output of the correlation indexed the quality of the attentional priority signal, with low (high) correlations reflecting weak (strong) feature selection. Separate correlations were calculated for correct and incorrect trials.

BOLD response patterns were composed of voxel-wise response amplitudes that were estimated using the deconvolution analysis above, and represented the spatial pattern of neural activity evoked when each feature was viewed or attended. Attention patterns were constructed using the response to CW and CCW-cued motion, contingent on whether the behavioral response was correct or incorrect during the attention task (four patterns in total). Baseline patterns were constructed using the response to correct rejections for CW and CCW motion during the baseline task (two patterns in total). Because both the speedup (target) and its detection were absent during correct rejections, the associated neural response should have reflected feature processing, without the contribution from neural activities related to target detection. Correlations of attention and baseline patterns with matching features (e.g., attention CW and baseline CW) were calculated. The resulting correlation coefficients were averaged across features and statistical inferences were conducted on Fisher-transformed values. For each ROI, we conducted planned t tests between the feature selection on correct and incorrect trials.

Multivariate analysis: multivoxel pattern classification For a separate measure of attentional priority, we trained a linear support vector machine (SVM; LIBSVM implementation; Chang and Lin, 2011) to discriminate between CW and CCW motion in the baseline task and then tested its ability to decode the motion directions when they were attended in the attention task, contingent on behavioral accuracy.

For this analysis, voxels were ranked by their r^2 value in the baseline task and the top 55 voxels in each ROI were used. Therefore, classification was based on the same number of voxels in each area. We note that our results were qualitatively identical when 35–145 voxels were used. We obtained single-trial BOLD responses (instances) for each voxel by averaging the time series between the time of peak response, as defined by a subject's response profile, to the shortest possible trial duration (8.8 s; 5th time point). Due to the variable time-to-peak across subjects, the average window contained between one and three time points. Each instance was treated as a point in 55-dimensional space and was used to populate multivoxel responses (classes), contingent on trial type. CW and CCW Baseline classes were composed

of BOLD response to correct rejection trials in the baseline task. CW and CCW Attention classes were composed of CW and CCW-cued trials in the attention task; separate classes were generated for correct and incorrect trials. Baseline classes were z-scored and used to train the SVM. Attention classes were z-scored using the mean and standard deviation of the training set, and used to test the SVM's accuracy in predicting the attended feature. Correct and incorrect trials were tested separately. Classification accuracy for each ROI was assessed as the average across hemispheres. Planned *t* tests were conducted in each ROI to compare classification accuracy between correct and incorrect trials.

For ROIs that exhibited a difference in classification accuracy between correct and incorrect trials, we conducted a permutation analysis to assess whether classification accuracy was significantly above or below chance level. For each subject, SVMs were trained with Baseline classes and tested with random samples from Attention classes. Specifically, trial labels for all four Attention classes (accuracy crossed with motion direction) were shuffled and split into two classes of equal size. Then, classification accuracy was assessed for the shuffled data. At the individual subject level, this process was repeated 10,000 times to create a null distribution of classification accuracy for the ROI. Null distributions were averaged across hemispheres. To construct the group-level null distribution, a single value was randomly selected from each subject's null distribution and the average value across subjects was calculated; this process was repeated 10,000 times to derive a null distribution for the ROI. Group-average classification accuracies below the 2.5th or above the 97.5th percentile were considered significantly below or above chance, respectively.

Note that the training and test data were based on different tasks in separate scanning runs and hence, entirely independent. Therefore, no leave-one-run-out cross-validation was necessary. On average, the training data contained 26 CW trials and 25 CCW trials. The test data contained 48 CW and 46 CCW correct trials, and 23 CW and 25 CCW incorrect trials.

Voxel exclusion criterion Voxels with responses larger than 5% signal change were excluded from all analyses as they presumably reflected noise. At the group level, this criterion removed 4.1% of voxels across ROIs. We note that the exact exclusion criterion did not qualitatively impact our results.

Visualization of group data

All analyses were performed on individual subject data with predefined ROIs and all quantitative results reported were based on averages across individual subject results. We also performed group averaging of the individual maps to provide a visualization of the overall pattern of brain activity during the attention and baseline tasks. Each subject's two hemispherical surfaces were first imported into Caret and affine-transformed into the 711-2B space of the Washington University at St. Louis (Buckner et al., 2004). The surface was then inflated to a sphere and six landmarks were drawn, which were used for spherical registration to the landmarks in the Population-Average,

Landmark- and Surface-based (PALS) atlas (Van Essen, 2005). Individual maps were transformed to the PALS atlas space and thresholded at a r^2 value of 0.12 in combination with a cluster constraint of 50 voxels. PALS-transformed maps were averaged across subjects and used solely for visualization purposes.

TMS experiments

Subjects

In total, 27 subjects participated in the TMS experiments, with 12 participating in each experiment. Out of the total, one participated in all four experiments (author M.G.), four participated in three experiments (L/R parietal and MT+), eight participated in two experiments (L/R parietal: 3; R parietal and MT+: 3; one author, M.J.); L parietal and MT+: 1; MT+ and sham: 1), and 13 participated in one experiment (L parietal: two subjects; MT+: 1; sham: 10). All subjects were neurologically intact, had normal or corrected-to-normal vision, and were recruited from the Michigan State University community (undergraduate and graduate students). They gave written informed consent under the study protocol approved by the Institutional Review Board at Michigan State University and were remunerated at a rate of \$15 per hour.

Repetitive TMS (rTMS): task and procedure

Each TMS experiment comprised of three sessions that were completed on separate days: one thresholding/practice session and two TMS sessions. A brain area was targeted unilaterally in each experiment.

We used the baseline and attention tasks from the fMRI experiment with the following modifications: (1) the staircase was specified to maintain a hit rate of 80%; (2) each run contained 100 trials; (3) to equate the duration of a block with that of the stimulation protocol (600 s), the dot fields were displayed for 3.7 s and the ITIs were changed to 1.35 s on 40% of trials, 1.8 s on 30% of trials, 2.25 s on 20% of trials, and 2.7 s on 10% of trials; (4) 70% of trials in the attention task were valid; (5) present and absent responses were always mapped to the index and middle finger, respectively; and (6) the central cross was always gray.

The thresholding/practice session occurred on a different day before the TMS session. During thresholding, subjects performed three to six runs of each task until performance was titrated to the expected level. The corresponding speedup magnitude served as the subject's threshold and was used during TMS sessions. Ultimately, two thresholds were obtained for each subject, one for the baseline task and one for the attention task.

Each subject performed both tasks (attention and baseline) across the two TMS sessions (i.e., one task per session). During each session, subjects performed two runs of a task, one before the TMS protocol (prestimulation; see *TMS protocol*) and another immediately after (poststimulation). Task order was counterbalanced across subjects. Task performance was evaluated as the difference between the hit rate and false alarm rate. For subjects that participated in two or more experiments, only one full thresholding session was conducted. During further TMS experiments, their thresholds were simply

updated by performing 50–100 thresholding trials immediately before their prestimulation run.

Sham TMS: task and procedure

The sham TMS experiment comprised of two sessions that were completed on separate days. In each session, subjects first performed 50–100 staircase trials to titrate performance to the expected level. The corresponding speedup magnitude served as their threshold and was used in subsequent blocks. Subjects then performed two blocks of the attention or baseline task separated by 10 min of sham TMS (see TMS protocol); task order was counterbalanced across subjects. We opted to combine the thresholding and stimulation sessions into a single session to exacerbate any potential fatigue effects.

TMS protocol

Three ROIs were targeted unilaterally for stimulation, left and right IPS1, and right MT+, in three separate experiments. IPS1 was topographically defined in individual subjects using the retinotopic mapping procedure while MT+ was defined using the MT+ localizer. In a given experiment, a ROI was overlaid on the corresponding T1-weighted anatomic MR image for each subject, and its centroid (visually determined) served as the target site. We used frameless stereotaxy (Brainsight 2, Rogue Research) in conjunction with a Polaris infrared positioning system (Northern Digital, Waterloo, Canada) to precisely place the coil over the target site; the coil was positioned with its handle $\sim 45^\circ$ from the midsagittal axis. During sham TMS, the coil was centered on the midsagittal plane but its face was rotated 90° away from the subject's head; thus, no cortical area received any direct stimulation.

Ten minutes of 1-Hz rTMS (600 pulses) was delivered to the target site using a Magstim Rapid² stimulator and a Magstim double 70 mm air film coil (The Magstim Company). Stimulation was delivered at a fixed intensity of 70% of maximum stimulator output. This stimulation intensity was chosen because several studies have found attentional effects when stimulating the parietal cortex at similar intensities (Hilgetag et al., 2001; Thut et al., 2005; Morris et al., 2007; Schenkluhn et al., 2008; Szczepanski and Kastner, 2013). We did not use motor threshold to determine stimulation intensity because it is not necessarily a reliable index of excitability in nonmotor cortical areas (Stewart et al., 2001; Robertson et al., 2003). For this reason, and to limit the length of the experiment and the total number of TMS pulses subjects received, motor threshold was not assessed.

Results

Behavior in the scanner: attention task

Subjects performed a speed detection task during which they were cued to selectively attend to one of two overlapping dot fields, one rotating CW and another rotating CCW (Fig. 1A). On each trial, a speedup (target) occurred in either the cued direction (valid, 80% of trials) or in the uncued direction (invalid, 20% of trials) and subjects were instructed to report whether or not they perceived the speedup in the cued direction. The magni-

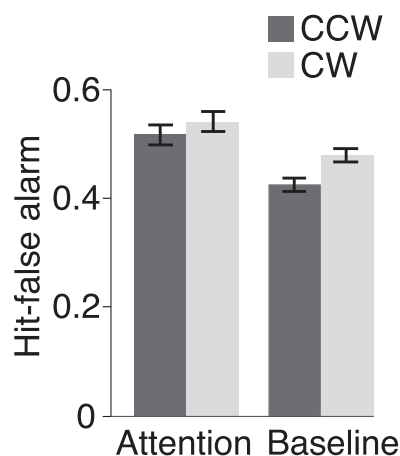


Figure 2. Behavioral results in the scanner. Error bars indicate \pm within-subject SEM following the method of Cousineau (2005).

tude of the speedup was controlled by an adaptive staircase that maintained the hit rate at 65%. This 4:1 validity ratio and intermediate performance level allowed a sufficient number of incorrect (miss) trials to be collected while ensuring that performance would benefit from valid cues. On average, we obtained 95 correct and 48 incorrect trials for each subject.

Consistent with our expectation, subjects successfully used the cue. The speed increment was detected near the expected hit rate ($\sim 65\%$) and the false alarm rate was low ($<15\%$). Additionally, the difference between hit and false alarm rates (i.e., hit-false alarm) revealed no difference in the ability to attend to either direction ($t_{(11)} = 0.65$, $p = 0.53^{a1}$; Fig. 2). All statistics are summarized in Table 1.

Behavior in the scanner: baseline task

In the same scanning session, subjects performed another speed detection task during which they were presented with a single dot field (Fig. 1B). On each trial, the dot field rotated either CW or CCW and a speedup occurred in 70% of trials. On the remaining 30% of trials, no speedup occurred. The magnitude of the speedup was controlled by an adaptive staircase that maintained the hit rate at 50%. This experimental design kept subjects engaged in the task and allowed us to obtain a sufficient number of correct rejection trials (when no speedup was presented and subjects correctly reported its absence) for subsequent analyses. On average, we obtained 51 correct rejection trials for each subject.

Behavioral performance indicated that subjects were successful at detecting the speedup in both directions. The hit rate was near the expected performance level and the false alarm rate was low ($<5\%$). The difference in performance during CW and CCW motion was marginally significant ($t_{(11)} = 2.18$, $p = 0.052^{a2}$), suggesting that it was easier to detect a CW target (Fig. 2). This difference was unexpected. However, we note that any performance difference in the baseline task should not affect our fMRI analyses because only correct rejection trials (i.e., trials without a target) were used.

Table 1. Statistics table

Line	Data/dependent variable	Type of test	Statistic	Confidence
	Results: behavior in scanner			
a1	Hit – false alarm	Paired <i>t</i> test (CW vs CCW)	$t = 0.65$; DoF = 11;	$p = 0.53$; CI = -0.10/0.057
a2		Paired <i>t</i> test (CW vs CCW)	$t = 2.18$; DoF = 11;	$p = 0.052$; CI = -0.11/0.0004
	Results: fMRI			
c1	BOLD response amplitude by condition	Two-way ANOVA (accuracy \times direction)	$F = 5.51$; DoF = (1,11)	$p = 0.039$; general $\eta^2 = 0.014$
c2		Two-way ANOVA (accuracy \times direction)	$F = 8.30$; DoF = (1,11)	$p = 0.015$; general $\eta^2 = 0.018$
d1	Correlation coefficients by accuracy condition (correct, incorrect)	Paired <i>t</i> test	$t = 3.21$; DoF = 11;	$p = 0.008$; CI = 0.031/0.17
d2		Paired <i>t</i> test	$t = 3.30$; DoF = 11;	$p = 0.007$; CI = 0.036/0.18
d3		Paired <i>t</i> test	$t = 3.71$; DoF = 11;	$p = 0.004$; CI = 0.061/0.24
d4		Paired <i>t</i> test	$t = 2.62$; DoF = 11;	$p = 0.024$; CI = 0.024/0.27
d5		Paired <i>t</i> test	$t = 2.60$; DoF = 11;	$p = 0.025$; CI = 0.016/0.19
d6		Paired <i>t</i> test	$t = 2.20$; DoF = 11;	$p = 0.049$; CI = 0.0002/0.20
e1	Classification accuracies by accuracy condition (correct, incorrect)	Paired <i>t</i> test	$t = 2.34$; DoF = 11;	$p = 0.039$; CI = 0.0028/0.093
e2		Permutation test (correct vs null)		$p = 10^{-3}$
e3		Permutation test (incorrect vs null)		$p = 10^{-4}$
	Results: TMS			
f1	Hit – false alarm	Paired <i>t</i> test (pre- vs poststimulation)	$t = 3.18$; DoF = 11;	$p = 0.009$; CI = 0.025/0.14
f2		Paired <i>t</i> test (pre- vs poststimulation)	$t = 0.12$; DoF = 11;	$p = 0.91$; CI = -0.069/0.076
f3		Paired <i>t</i> test (pre- vs poststimulation)	$t = 3.36$; DoF = 11;	$p = 0.006$; CI = 0.034/0.16
f4		Paired <i>t</i> test (pre- vs poststimulation)	$t = 0.57$; DoF = 11;	$p = 0.58$; CI = -0.044/0.075
f5		Two-way ANOVA (task \times stimulation period)	$F = 9.40$; DoF = (1,11)	$p = 0.011$; general $\eta^2 = 0.018$
f6		Two-way ANOVA (task \times stimulation period)	$F = 5.28$; DoF = (1,11)	$p = 0.042$; general $\eta^2 = 0.010$
f7		Paired <i>t</i> test (pre- vs poststimulation)	$t = 2.82$; DoF = 11;	$p = 0.017$; CI = 0.021/0.17
f8		Paired <i>t</i> test (pre- vs poststimulation)	$t = 2.26$; DoF = 11;	$p = 0.045$; CI = 0.0016/0.11
f9		Two-way ANOVA (task \times stimulation period)	$F = 11.62$; DoF = (1,11)	$p = 0.006$; general $\eta^2 = 0.13$
f10		Two-way ANOVA (task \times stimulation period)	$F = 0.90$; DoF = (1,11)	$p = 0.36$; general $\eta^2 = 0.009$
f11		Paired <i>t</i> test (pre- vs poststimulation)	$t = 1.60$; DoF = 11;	$p = 0.14$; CI = -0.14/0.021
f12		Paired <i>t</i> test (pre- vs poststimulation)	$t = 0.34$; DoF = 11;	$p = 0.74$; CI = -0.073/0.053

Attention and baseline tasks modulate BOLD response in occipital and frontoparietal areas

To identify cortical areas that were modulated by the tasks, we performed a subject-based deconvolution analysis. Voxel-wise deconvolved responses were computed for each condition and the amount of variance in the time course accounted for by the deconvolution model (r^2) represented the extent of task modulation. Group-averaged r^2 maps were visualized in the PALS atlas space using spherical registration (for details, see Materials and Methods).

At the group-level, a network of frontoparietal areas, as well as occipital visual areas, showed significant modulation by the both the attention (Fig. 3A) and baseline (data not shown) tasks. This overall pattern of activation was similar to findings from many previous studies of attentional control (Kastner and Ungerleider, 2000; Corbetta and Shulman, 2002). Active areas in the occipital and parietal cortices overlapped with areas defined via retinotopy. Additionally, we defined two frontal ROIs (FEF and IFJ) in each subject using their r^2 map of the attention task. FEF was defined near the junction of the precentral and superior frontal sulcus and IFJ was defined near the

junction of the precentral and inferior frontal sulcus. These two regions appeared contiguous on the group-averaged map, but they formed distinct clusters in individual subject maps. In each subject, we obtained 11 ROIs: V1, V2, V3, V3A/B, V4, MT+, V7/IPS0, IPS1, IPS2, FEF, and IFJ. No laterality was observed in the results for each area, therefore we averaged the data from corresponding areas across the two hemispheres.

BOLD amplitude does not vary with task performance

We observed robust task-related BOLD responses in all of our ROIs and, for illustrative purposes, we plotted the response of V1 and IPS1 (Fig. 3B). The BOLD response to all conditions in the attention task (accuracy crossed with cued direction, CW and CCW) peaked between 4.4 and 8.8 s after trial onset (time points 3–5) and was very similar between conditions. To quantify this observation, we computed the average response amplitude for each condition and conducted a two-way repeated-measures ANOVA (two accuracy conditions \times two attended directions). The analysis revealed a main effect of direction in

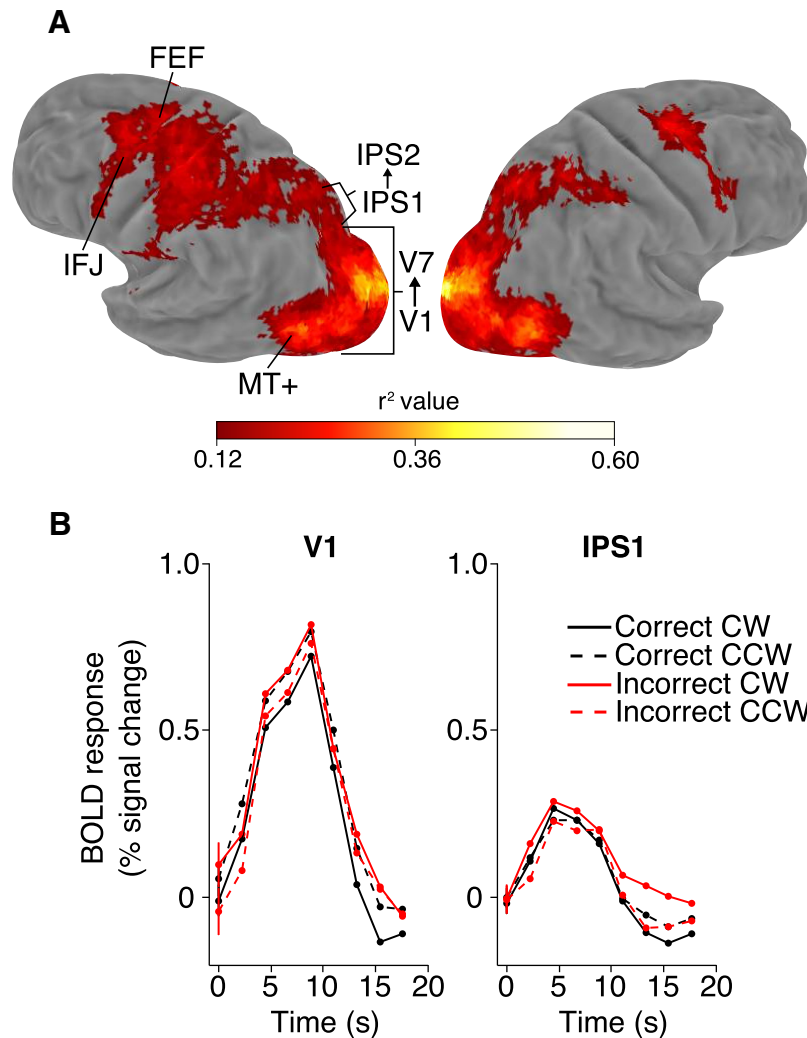


Figure 3. Univariate results. **A**, Group r^2 map of the attention task shown on an inflated Caret atlas surface. The approximate locations of retinotopically-defined (V1-7, MT+, IPS1, and IPS2) and task-defined (FEF and IFJ) areas are indicated by lines. **B**, Mean BOLD response in the attention task from two ROIs (V1 and IPS1). The error bar on the first time point is the average \pm within-subject SEM across all time points.

IPS2, response amplitude was larger for CW than CCW attention ($F_{(1,11)} = 5.51, p = 0.039^{c1}$), and an interaction in V1, when correct, response amplitude was larger for CCW than CW attention; when incorrect, the pattern was reversed ($F_{(1,11)} = 8.30, p = 0.015^{c2}$). Importantly, the main effect of accuracy was nonsignificant in all ROIs (all $p > 0.2$), therefore, overall BOLD amplitude did not vary with performance. This result is thus inconsistent with an interpretation that performance variation in the attention task was caused by fluctuations of general behavioral state (e.g., fatigue or vigilance), as such effects are known to be reflected in overall BOLD amplitude variations (Boly et al., 2007; Esterman et al., 2013).

Average neural patterns for prioritized features in occipital and frontoparietal areas vary with task performance

To examine a possible relationship between attentional priority and task performance, we used a correlation analysis to compare the spatial pattern of neural activity when

subjects were correct versus incorrect in the attention task (Fig. 4). Voxel-wise BOLD responses were used to obtain spatial patterns of neural activity for attended and isolated features (CW and CCW motion). Separate Attention patterns were constructed for CW and CCW-cued trials in the attention task, and baseline patterns were constructed with voxel responses during correct rejection trials in the baseline task. The response to correct rejections was used because there was neither a physical nor perceived target during these trials, which made them suitable for isolating feature-specific neural responses, uncontaminated by target-related responses. We used the correlation between baseline and attention patterns to index the quality of attentional priority, with low (high) correlations reflecting weak (strong) feature selection.

We conducted planned comparisons between the quality of attentional priority on correct and incorrect trials within each ROI (Fig. 5). Because all visual areas showed similar results and because we are primarily interested in frontoparietal areas' role in attentional control, we aggre-

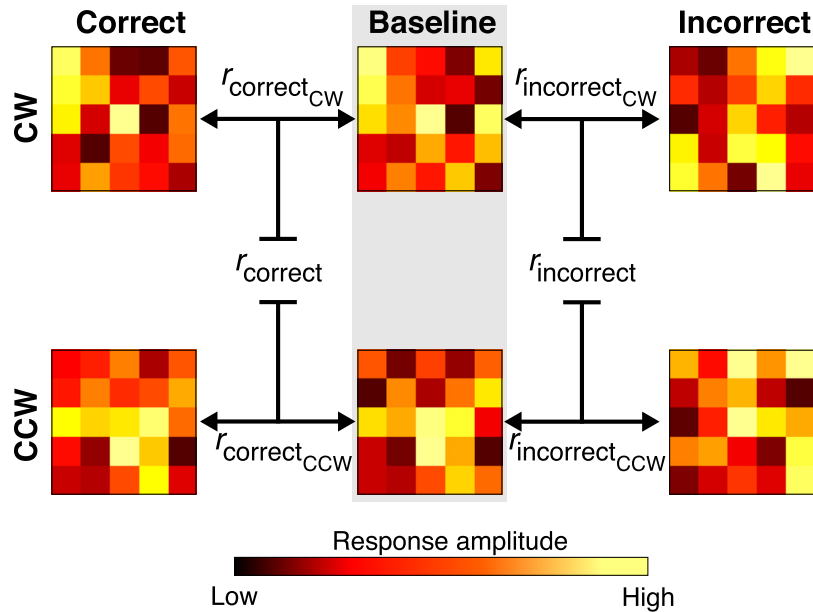


Figure 4. Schematic of the correlation analysis. Each matrix represents the spatial pattern of response amplitudes from voxels within a ROI (amplitude is color coded according to the scale bar at the bottom). The middle column of matrices (shaded area labeled “baseline”) illustrates the baseline neural response pattern to each direction (CW and CCW) during the baseline task. The other two columns illustrate the neural response pattern to an attended motion direction during correct and incorrect trials during the attention task. The black double-sided arrows between matrices represent the correlations that were calculated (Pearson’s r) and the bounded lines represent the averaging of correlation coefficients across directions to obtain an overall index of attentional priority quality for correct and incorrect trials.

gated data from extrastriate visual areas (abbreviated to ExS in figures) by averaging the correlation coefficients across V2, V3, V3A/B, and V4. We kept MT+ separate as there is a strong a priori link between MT+ activity and motion processing. We found that Baseline and Attention patterns were more correlated for correct than incorrect trials in V1 ($t_{(11)} = 3.21, p = 0.008^{d1}$), extrastriate regions ($t_{(11)} = 3.30, p = 0.007^{d2}$), MT+ ($t_{(11)} = 3.71, p = 0.004^{d3}$), V7/IPS0 ($t_{(11)} = 2.62, p = 0.024^{d4}$), IPS1 ($t_{(11)} = 2.60, p = 0.025^{d5}$), and IFJ ($t_{(11)} = 2.20, p = 0.049^{d6}$). These results demonstrate that the spatial pattern of neural activity within these cortical areas encodes a more veridical representation of the attended feature during correct trials. The observed neural-behavioral correlates in posterior parietal and inferior frontal areas support their role in encoding attentional priority for features, whereas the analogous effects in visual areas likely reflect attentional modulation due to feedback. In the following, we sought converging evidence with a classification approach that provided another measure of pattern similarity.

Feature coding in posterior parietal cortex (PPC) tracks trial-by-trial fluctuations in task performance

The correlation analysis above used the average neural response across trials; yet, attentional control likely fluctuates across individual trials, leading to different behavioral outcomes. Therefore, a trial-by-trial assessment would reveal cortical areas that consistently encode the attended feature and facilitate task performance.

A classification approach was used to assess feature representation on individual trials. A linear classifier was trained with correct rejection trials from the baseline task

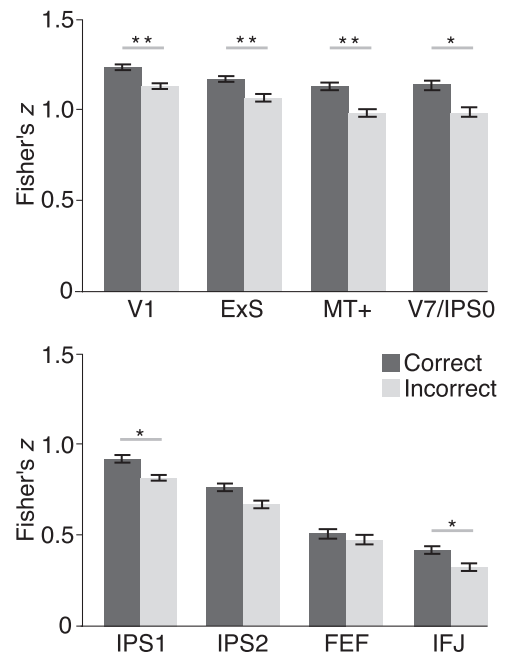


Figure 5. Average neural patterns for prioritized features in occipital and frontoparietal areas vary with task performance. Group-average Fisher-transformed correlation coefficients (averaged across motion directions) are shown, which reflect the similarity of neural patterns of activity between the attention and baseline tasks for correct and incorrect trials. The ExS label represents extrastriate visual areas. Error bars are \pm within-subject SEM following the method of Cousineau (2005). Asterisks indicate the significance level in paired t tests (** $p < 0.01, *p < 0.05$).

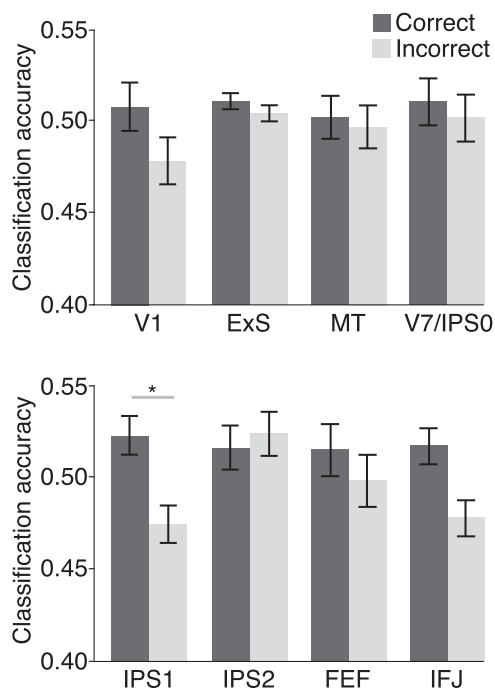


Figure 6. Feature coding in PPC tracks trial-by-trial fluctuations in task performance. Group-average classification accuracies are shown and plotting conventions are the same as Figure 5.

to discriminate between CW and CCW motion. Then, the classifier was tested on how well it decoded the attended direction in individual correct and incorrect trials of the attention task. For each ROI, classification accuracy for correct and incorrect trials were compared (Fig. 6), revealing that the attended direction was better discriminated for correct than incorrect trials in IPS1 ($t_{(11)} = 2.34$, $p = 0.039^{e1}$). Moreover, by constructing a null-distribution that characterized chance-level classification performance, we found that the attended feature was discriminated above chance for correct trials ($p = 10^{-3}$) and below chance for incorrect trials ($p = 10^{-4}$). These results indicate that IPS1 consistently encodes a more veridical representation of the attended feature during correct than incorrect trials. The below-chance classification for incorrect trials is somewhat unexpected; this could suggest that errors were partly due to subjects attending to the uncued feature on those trials, which would lead to opposite neural response patterns for the training and test datasets. Overall, both our correlation and classification approaches provide converging evidence of IPS1's role in the maintenance of attentional priority.

PPC is necessary for feature-based attention

Because our multivariate analyses indicated that neural activity patterns in IPS1 correlate with feature selection and task performance, we sought to assess whether it is causally involved in feature-based attention. Subjects performed the attention and baseline tasks while neuro-navigated rTMS was used to disrupt neural processing unilaterally in left or right IPS1. Stimulation was centered on each retinotopically-defined ROI, but given the rela-

tively small size of these cortical areas compared to the estimated spatial extent of TMS stimulation (Walsh and Pascual-Leone, 2003), we refer to the area stimulated as the PPC.

We reasoned that because feature-based selection was necessary in the attention task, but not the baseline task, disruption of areas primarily representing attentional priority would impair performance in the attention task and leave performance in the baseline task largely unchanged. With this rationale, we conducted two separate experiments in which rTMS was applied unilaterally to left or right PPC. Before TMS sessions, speedup magnitude thresholds for the attention and baseline tasks were obtained for each subject such that hit rates were equated. Then, on a separate day, subjects performed two blocks of either task, separated by 1-Hz rTMS. After stimulation of left PPC (Fig. 7A), performance was significantly impaired in the attention task ($t_{(11)} = 3.18$, $p = 0.009^{f1}$) but not the baseline task ($t_{(11)} = 0.12$, $p = 0.91^{f2}$). Similarly, stimulation of right PPC (Fig. 7B) impaired performance in the attention task ($t_{(11)} = 3.36$, $p = 0.006^{f3}$) but not the baseline task ($t_{(11)} = 0.57$, $p = 0.58^{f4}$). Consistent with these planned comparisons, a two-way repeated measures ANOVA (two tasks \times two stimulation periods) revealed a significant interaction effect for left ($F_{(1,11)} = 9.40$, $p = 0.011^{f5}$) and right ($F_{(1,11)} = 5.28$, $p = 0.042^{f6}$) parietal stimulation sites. Thus, these results support a causal role of PPC in the maintenance of attentional priority that mediates feature selection.

Neurodisruption of MT+ does not selectively impair feature-based attention

Because our fMRI correlation analysis also revealed performance correlates in visual areas, we conducted a third TMS experiment where we stimulated right MT+ to further dissociate visual areas from control areas of attention (Fig. 7C). Given its critical role in visual motion processing (Newsome and Paré, 1988), we reasoned that neurodisruption of MT+ should lead to equivalent decrements in performance for the attention and baseline tasks because both rely on basic motion processing. Consistent with this prediction, rTMS impaired performance in the attention ($t_{(11)} = 2.82$, $p = 0.017^{f7}$) and baseline task ($t_{(11)} = 2.26$, $p = 0.045^{f8}$). This was verified by a two-way repeated measures ANOVA that revealed a main effect of stimulation period ($F_{(1,11)} = 11.62$, $p = 0.006^{f9}$) and no interaction effect ($F_{(1,11)} = 0.90$, $p = 0.36^{f10}$).

General behavioral state does not explain performance decrements due to rTMS

A potential explanation for the performance decrements observed in the above TMS experiments might relate to variations in general behavioral state (e.g., fatigue or vigilance). In particular, impaired performance post stimulation could have arisen simply because subjects were fatigued after completing their first block of trials, and this fatigue might be particularly severe for the attention task. To assess this possibility, we performed a fourth experiment in which no cortical area was directly stimulated. Instead, the coil face was oriented 90° away from the scalp (sham TMS). In addition, we exacerbated po-

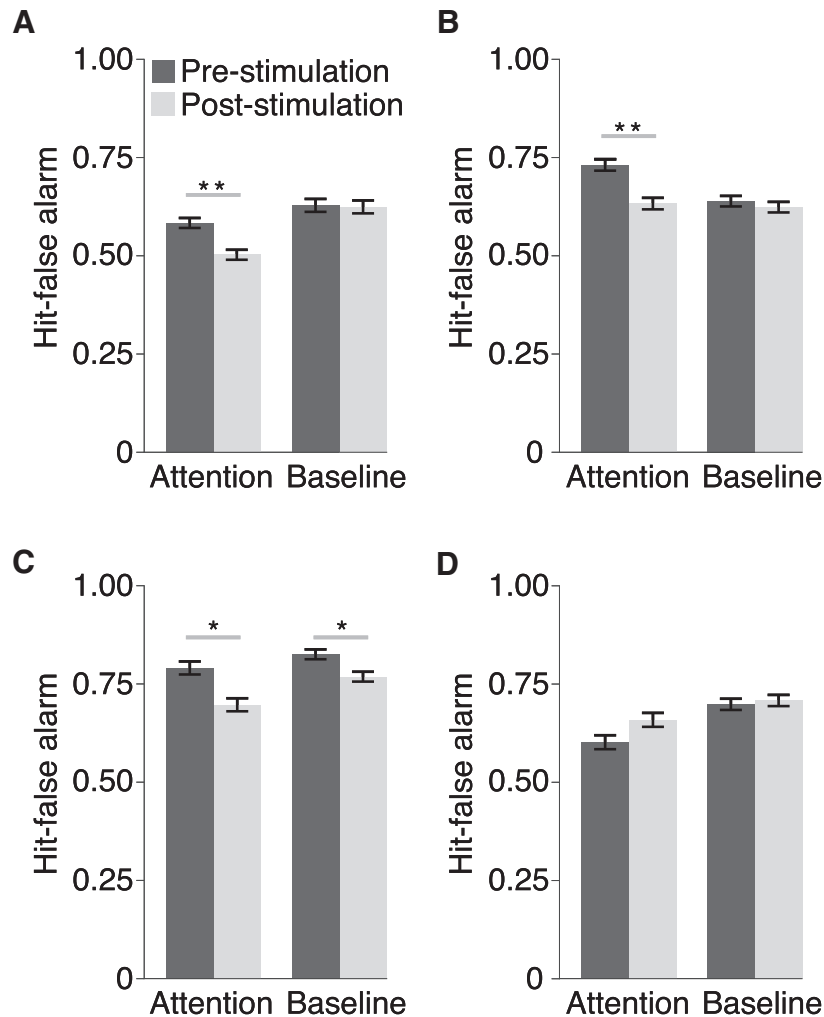


Figure 7. Results of the TMS experiments. Each panel presents the results for stimulation centered on (A) left IPS1, (B) right IPS1, (C) right MT+, and (D) sham TMS. Plotting conventions are the same as Figure 5.

tential fatigue effects by combining thresholding and TMS sessions into a single session. Thus, subjects first completed one to two blocks of staircase trials, and then immediately performed the pre- and post-stimulation blocks, separated by sham TMS. We reasoned that if fatigue contributed to our TMS effects, then performance would be impaired in the post-block even without brain stimulation. Our results reject this hypothesis (Fig. 7D): performance was unchanged between pre- and post-blocks in the attention task ($t_{(11)} = 1.60, p = 0.14^{f11}$) and baseline task ($t_{(11)} = 0.34, p = 0.74^{f12}$). Thus, fatigue cannot account for the observed effects of PPC and MT+ stimulation.

Overall, we found that applying rTMS to PPC produced a selective deficit in the attention task, demonstrating its specific role in attentional selection. This was in contrast to MT+ stimulation that produced a deficit in both tasks, demonstrating its general role in motion processing. Thus, our results provide converging evidence that PPC, and IPS1 in particular, is causally involved in the maintenance of attentional priority.

Discussion

In this study, we examined how neural activity in dorsal frontoparietal areas is related to behavioral performance in a feature-based attention task. We found that attending to a feature (motion direction) produced spatial patterns of neural activity in frontal, parietal, and visual areas that better resembled those of an isolated feature (evoked by a single motion direction) during correct than during incorrect trials. On a trial-by-trial basis, this pattern congruency effect was uniquely found in IPS1, as revealed by pattern classification analyses. Additionally, we found that rTMS centered on IPS1 led to a performance impairment in the attention but not the baseline task, whereas rTMS to MT+ led to equivalent impairments in both tasks. Finally, sham TMS did not change performance in either task, ruling out the influence of general behavioral state on task performance. These results reveal that frontoparietal areas maintain attentional priority that facilitates the selection of visual features, and in particular, PPC, including IPS1 plays a causal role in such selection.

Previous studies have provided indirect evidence that frontoparietal neural activity is related to the control of feature-based attention (Serences and Boynton, 2007; Liu et al., 2011; Ester et al., 2016). However, whether these neural signals mediate task performance and how such signals represent the attended feature are unclear. Here, we addressed both issues by examining the relationship between patterned neural activity and task performance, as well as the consequence of neural perturbation on task performance.

The baseline task provided a measure of the neural pattern evoked by each motion direction without selective attention. To the extent that subjects successfully attended to the cued direction in the attention task, the observed neural pattern should resemble the baseline pattern for that direction and facilitate performance in a difficult threshold-level detection task. Indeed, we found such an effect in visual areas including V1 and frontoparietal areas (V7/IPS0, IPS1, and IFJ). In these areas, the correlation between the attention pattern and the baseline pattern was greater for correct than incorrect trials, suggesting that feature-based attention operates by biasing the population activity toward that evoked by the feature alone. These results are consistent with the finding that neuronal tuning in V4 is shifted to the attended feature (David et al., 2008) as well as with the finding that attention shifts fMRI voxel's semantic category representations during visual search in natural movies (Çukur et al., 2013). Our results are complementary to these previous findings and go beyond them by demonstrating that attention-induced shifts in neural population activity are functionally significant in that they correlate with task performance.

We also used a pattern classification approach to assess the relationship between neural activity patterns and task performance. Because the classification approach uses single-trial data, this allows us to examine how trial-by-trial variations in neural activity contribute to attentional selection. This analysis showed that the pattern difference between the two features, the information extracted by the classifier, was more aligned between the baseline and the attention tasks on correct trials. This result suggests that to the extent that the brain could rely on the same discriminative information between the two features when they were presented alone, successful attentional selection can be achieved when they were presented together in competition. We note that we only found this result in IPS1, while the correlation analysis found neural correlates of behavioral accuracy in multiple cortical areas. The difference in results could be due, in part, to the use of mean patterns across trials in the correlation analysis that reflect an overall effect of attention, and the use of single-trial BOLD responses in the classification analysis, which are sensitive to trial-level fluctuations. Although, in principle, such fluctuations should manifest both at the source and the destination of attentional modulation, it is plausible that neural conduction would introduce additional noise at the destination relative to the source. If so, it would be easier to detect a performance-based effect at the source region than the destination region. Thus, our classification results indi-

rectly support the notion that IPS1 contains the source of attentional control that modulates visual areas during feature-based selection. In addition, single-trial neural patterns used in the classification analysis are likely to be quite noisy, which could make it difficult to achieve reliable pattern separation. Indeed, the overall classification accuracy in our results was rather low, presumably limited by the noisy estimate of the trial-level responses and the limited size of the training dataset. Although the effect in classification accuracy was somewhat weak, it provided additional evidence that IPS1 is an important cortical area in shaping attentional selection, and guided our selection of PPC as a stimulation target during TMS.

To obtain converging evidence, we used rTMS to disrupt local neural processing in PPC to test its causal role in feature-based attention. We reasoned that if PPC is causally involved in attentional selection, neurodisruption should produce behavioral deficits in the attention task and not the baseline task. Importantly, performance in both tasks was titrated to be at equivalent levels with adaptive staircase methods. Therefore, general task difficulty cannot explain any differential effect produced by neurodisruption. Our results support this prediction: disrupting PPC produced a behavioral deficit in the attention task but not in the baseline task. We dissociated this selective top-down attentional impairment from bottom-up motion processing by stimulating MT+, a region with strong a priori links to bottom-up processing of visual motion. Because each task relies on basic motion processing, we reasoned that neurodisruption should produce equivalent behavioral deficits in both tasks. Consistent with this prediction, disrupting MT+ yielded equivalent behavioral deficits in both tasks. Finally, we dissociated the effects of neurodisruption from variations in general behavioral state (e.g., fatigue or vigilance) with sham TMS and protracted experimental sessions. We reasoned that if variations in task performance are due to fatigue, behavioral deficits should be observed in both tasks even without neurodisruption. Our results reject this hypothesis as performance was unchanged with sham TMS. Overall, these findings demonstrate that neural activity in PPC, and in particular, IPS1, is causally and specifically involved in the control of feature-based attention.

Our TMS results apparently contradict those of a previous study that found that stimulation of anterior, but not posterior, IPS impaired performance in a feature-cued visual search task (Schenkluhn et al., 2008). However, there are some critical differences between the TMS protocol used in the two studies. They used an online protocol to deliver brief TMS pulses right after the cue, which was presumed to only affect neural activity during the cue period. In contrast, the effects of our offline protocol have been shown to persist during the test block (Hilgetag et al., 2001; Thut et al., 2005; Zanto et al., 2011). Therefore, it is possible that anterior IPS plays a role in cue processing (e.g., setting the task goal), whereas posterior IPS plays a critical role in actively maintaining the selection of a goal-relevant feature during stimulus presentation. Future studies could determine the time course of relevance for these areas by using TMS to disrupt anterior

and posterior IPS at different times during cue and stimulus presentation.

Our converging fMRI and TMS results highlight the role of the PPC in controlling feature-based attention by demonstrating a strong behavioral correlate between parietal activity and task performance and its causal role in guiding behavior. Considering the strong support of this area in the control of spatial attention (Bisley and Goldberg, 2010; Buschman and Kastner, 2015), the overall results suggest that parietal areas contain domain-general attentional control signals for both spatial and nonspatial attention. Our fMRI correlation results also revealed a neural correlate of task performance in IFJ. IFJ has been suggested to be a shared node between the dorsal and ventral attentional control network (Corbetta et al., 2008) and recent studies have demonstrated its role in object-based attention and working memory (Zanto et al., 2011; Baldauf and Desimone, 2014), which are closely related to feature-based attention. The overall results suggest that IFJ and PPC represent attentional priority of features by maintaining population neural activity similar to that evoked by those features alone. These neural signals could serve as attentional templates, which have been proposed to control the deployment of feature-based attention in theoretical models (Wolfe, 1994; Desimone and Duncan, 1995). We note, however, that the overall magnitude of correlation in IFJ is low and it also did not exhibit performance-related effects in the classification analysis. Therefore, its role in feature-based selection might be weaker than object-based selection or working memory. In addition, other areas in the frontoparietal network did not exhibit any fMRI correlate of performance. In particular, fMRI pattern similarity in IPS2 and FEF did not vary with task performance although previous studies have shown that neural signals in these areas can be used to decode the attended feature (Liu et al., 2011; Liu and Hou, 2013; Ester et al., 2016). It is conceivable that other types of control signals, the ones that do not necessarily resemble those that process the original features, also participate in attentional selection. For example, these areas in the network could encode more abstract, perhaps categorical, information that guides attentional selection. The precise role of individual cortical areas in coordinating attentional selection awaits further investigation.

More broadly, our results are also informative regarding the general role of PPC in visual processing. Traditionally, parietal cortex has been associated with visuospatial and visuomotor processing (Mishkin and Ungerleider, 1982; Andersen and Buneo, 2002). However, more recent work has demonstrated neural selectivity to many nonspatial properties in this cortical area, such as simple features (Sereno and Maunsell, 1998; Toth and Assad, 2002), arbitrary categories (Freedman and Assad, 2006; Fitzgerald et al., 2011), and even abstract identity information (Jeong and Xu, 2016). Our results are thus consistent with this emerging view of nonspatial representation in parietal cortex and further suggest that such nonspatial representations can facilitate the selection of behaviorally relevant visual features.

References

- Andersen RA, Buneo CA (2002) Intentional maps in posterior parietal cortex. *Annu Rev Neurosci* 25:189–220. [CrossRef Medline](#)
- Baldauf D, Desimone R (2014) Neural mechanisms of object-based attention. *Science* 344:424–427. [CrossRef Medline](#)
- Bisley JW, Goldberg ME (2010) Attention, intention, and priority in the parietal lobe. *Annu Rev Neurosci* 33:1–21. [CrossRef Medline](#)
- Boly M, Baeteau E, Schnakers C, Degueldre C, Moonen G, Luxen A, Phillips C, Peigneux P, Maquet P, Laureys S (2007) Baseline brain activity fluctuations predict somatosensory perception in humans. *Proc Natl Acad Sci USA* 104:12187–12192. [CrossRef Medline](#)
- Buckner RL, Head D, Parker J, Fotenos AF, Marcus D, Morris JC, Snyder AZ (2004) A unified approach for morphometric and functional data analysis in young, old, and demented adults using automated atlas-based head size normalization: reliability and validation against manual measurement of total intracranial volume. *Neuroimage* 23:724–738. [CrossRef Medline](#)
- Buschman TJ, Kastner S (2015) From behavior to neural dynamics: an integrated theory of attention. *Neuron* 88:127–144. [CrossRef Medline](#)
- Chang CC, Lin CJ (2011) LIBSVM: a library for support vector machines. *ACM Trans Intell Syst Technol* 2:27. [CrossRef](#)
- Corbetta M, Shulman GL (2002) Control of goal-directed and stimulus-driven attention in the brain. *Nat Rev Neurosci* 3:201–215. [CrossRef](#)
- Corbetta M, Patel G, Shulman GL (2008) The reorienting system of the human brain: from environment to theory of mind. *Neuron* 58:306–324. [CrossRef Medline](#)
- Cousineau D (2005) Confidence intervals in within-subject designs: a simpler solution to Loftus and Masson's method. *Tutor Quant Methods Psychol* 1:42–45. [CrossRef](#)
- Çukur T, Nishimoto S, Huth AG, Gallant JL (2013) Attention during natural vision warps semantic representation across the human brain. *Nat Neurosci* 16:763–770. [CrossRef Medline](#)
- David SV, Hayden BY, Mazer JA, Gallant JL (2008) Attention to stimulus features shifts spectral tuning of V4 neurons during natural vision. *Neuron* 59:509–521. [CrossRef Medline](#)
- Deco G, Rolls ET (2004) A neurodynamical cortical model of visual attention and invariant object recognition. *Vision Res* 44:621–642. [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 USA* 93:2382–2386. [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. [Medline](#)
- Ester EF, Sutterer DW, Serences JT, Awh E (2016) Feature-selective attentional modulations in human frontoparietal cortex. *J Neurosci* 36:8188–8199. [CrossRef Medline](#)
- Esterman M, Noonan SK, Rosenberg M, DeGutis J (2013) In the zone or zoning out? Tracking behavioral and neural fluctuations during sustained attention. *Cereb Cortex* 23:2712–2723. [CrossRef Medline](#)
- Fitzgerald JK, Freedman DJ, Assad JA (2011) Generalized associative representations in parietal cortex. *Nat Neurosci* 14:1075–1079. [CrossRef Medline](#)
- Freedman DJ, Assad JA (2006) Experience-dependent representation of visual categories in parietal cortex. *Nature* 443:85–88. [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](#)
- Hilgetag CC, Théoret H, Pascual-Leone A (2001) Enhanced visual spatial attention ipsilateral to rTMS-induced “virtual lesions” of human parietal cortex. *Nat Neurosci* 4:953–957. [CrossRef](#)

- Jeong SK, Xu Y (2016) Behaviorally relevant abstract object identity representation in the human parietal cortex. *J Neurosci* 36:1607–1619. [CrossRef](#) [Medline](#)
- Kastner S, Ungerleider LG (2000) Mechanisms of visual attention in the human cortex. *Annu Rev Neurosci* 23:315–341. [CrossRef](#) [Medline](#)
- Konen CS, Kastner S (2008) Representation of eye movements and stimulus motion in topographically organized areas of human posterior parietal cortex. *J Neurosci* 28:8361–8375. [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](#)
- Liu T, Hou Y (2013) A hierarchy of attentional priority signals in human frontoparietal cortex. *J Neurosci* 33:16606–16616. [CrossRef](#) [Medline](#)
- Mishkin M, Ungerleider LG (1982) Contribution of striate inputs to the visuospatial functions of parieto-preoccipital cortex in monkeys. *Behav Brain Res* 6:57–77. [Medline](#)
- Morris AP, Chambers CD, Mattingley JB (2007) Parietal stimulation destabilizes spatial updating across saccadic eye movements. *Proc Natl Acad Sci USA* 104:9069–9074. [CrossRef](#) [Medline](#)
- Nestares O, Heeger DJ (2000) Robust multiresolution alignment of MRI brain volumes. *Magn Reson Med* 43:705–715. [Medline](#)
- Newsome WT, Paré EB (1988) A selective impairment of motion perception following lesions of the middle temporal visual area (MT). *J Neurosci* 8:2201–2211. [Medline](#)
- Prins N, Kingdom FAA (2009) Palamedes: Matlab routines for analyzing psychophysical data. [Medline](#)
- Robertson EM, Théoret H, Pascual-Leone A (2003) Studies in cognition: the problems solved and created by transcranial magnetic stimulation. *J Cogn Neurosci* 15:948–960. [CrossRef](#) [Medline](#)
- Schenkluhn B, Ruff CC, Heinen K, Chambers CD (2008) Parietal stimulation decouples spatial and feature-based attention. *J Neurosci* 28:11106–11110. [CrossRef](#) [Medline](#)
- Schluppeck D, Glimcher P, Heeger DJ (2005) Topographic organization for delayed saccades in human posterior parietal cortex. *J Neurophysiol* 94:1372–1384. [CrossRef](#) [Medline](#)
- Serences JT, Boynton GM (2007) Feature-based attentional modulations in the absence of direct visual stimulation. *Neuron* 55:301–312. [CrossRef](#) [Medline](#)
- Sereno AB, Maunsell JH (1998) Shape selectivity in primate lateral intraparietal cortex. *Nature* 395:500–503. [CrossRef](#) [Medline](#)
- Sereno MI, Dale AM, Reppas JB, Kwong KK (1995) Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Science* 268:889–893. [Medline](#)
- Sereno MI, Pitzalis S, Martinez A (2001) Mapping of contralateral space in retinotopic coordinates by a parietal cortical area in humans. *Science* 294:1350–1354. [CrossRef](#) [Medline](#)
- Silver MA, Kastner S (2009) Topographic maps in human frontal and parietal cortex. *Trends Cogn Sci* 13:488–495. [CrossRef](#) [Medline](#)
- Simons DJ, Rensink RA (2005) Change blindness: past, present, and future. *Trends Cogn Sci* 9:16–20. [CrossRef](#) [Medline](#)
- Stewart LM, Walsh V, Rothwell JC (2001) Motor and phosphene thresholds: a transcranial magnetic stimulation correlation study. *Neuropsychologia* 39:415–419. [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](#)
- Szczepanski SM, Kastner S (2013) Shifting attentional priorities: control of spatial attention through hemispheric competition. *J Neurosci* 33:5411–5421. [CrossRef](#) [Medline](#)
- Thut G, Nietzel A, Pascual-Leone A (2005) Dorsal posterior parietal rTMS affects voluntary orienting of visuospatial attention. *Cereb Cortex* 15:628–638. [CrossRef](#) [Medline](#)
- Toth LJ, Assad JA (2002) Dynamic coding of behaviourally relevant stimuli in parietal cortex. *Nature* 415:165–168. [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](#)
- Walsh V, Pascual-Leone A (2003) *Transcranial magnetic stimulation: a neurochromometrics of mind*. Cambridge: MIT Press.
- 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. [Medline](#)
- Wolfe JM (1994) Guided search 2.0 a revised model of visual search. *Psychon Bull Rev* 1:202–238. [CrossRef](#) [Medline](#)
- Yantis S (2000) Goal-directed and stimulus-driven determinants of attentional control. *Atten Perform* 18:73–103.
- Zanto TP, Rubens MT, Thangavel A, Gazzaley A (2011) Causal role of the prefrontal cortex in top-down modulation of visual processing and working memory. *Nat Neurosci* 14:656–661. [CrossRef](#) [Medline](#)