

Attention modulates spatial priority maps in the human occipital, parietal and frontal cortices

Thomas C Sprague¹ & John T Serences^{1,2}

Computational theories propose that attention modulates the topographical landscape of spatial ‘priority’ maps in regions of the visual cortex so that the location of an important object is associated with higher activation levels. Although studies of single-unit recordings have demonstrated attention-related increases in the gain of neural responses and changes in the size of spatial receptive fields, the net effect of these modulations on the topography of region-level priority maps has not been investigated. Here we used functional magnetic resonance imaging and a multivariate encoding model to reconstruct spatial representations of attended and ignored stimuli using activation patterns across entire visual areas. These reconstructed spatial representations reveal the influence of attention on the amplitude and size of stimulus representations within putative priority maps across the visual hierarchy. Our results suggest that attention increases the amplitude of stimulus representations in these spatial maps, particularly in higher visual areas, but does not substantively change their size.

Prominent computational theories of selective attention posit that basic properties of visual stimuli are encoded in a series of interacting priority maps that are found at each stage of the visual system^{1–6}. The maps in different areas are thought to encode different stimulus features (for example, orientation, color or motion) on the basis of the selectivity of component neurons. Two general themes governing the organization of these maps have emerged. First, accurately encoding the spatial location of relevant stimuli is the fundamental goal of these priority maps, as spatial position is necessary to guide saccadic eye movements (and other exploratory and reflexive motor responses). Second, priority maps early in the visual system reflect primarily the physical salience of stimuli in the visual field, whereas priority maps in later areas increasingly index the behavioral relevance of stimuli independent of physical salience^{4,5}.

Although many studies have investigated the influence of spatial attention on single-unit neural activity over the last several decades^{7–17}, directly examining the impact of attention on the topographic profile across an entire spatial priority map is a major challenge because single units have access to a limited window of the spatial scene⁵. This is a key limitation, as the relationship between changes in the size and amplitude of individual spatial receptive fields (or voxel-level receptive fields across populations of neurons) and changes in the fidelity of population-level spatial encoding are not related in a straightforward manner (ref. 18 discusses this issue with respect to population codes for orientation). For example, if spatial receptive fields are uniformly shrunk by attention while viewing a stimulus, the population-level spatial representation (or priority map) carried by all of those neurons might shrink or become sharper, but the code may be more vulnerable to uncorrelated noise (as there is less redundant coding of any given spatial position by the population). Alternatively, a uniform increase in spatial receptive field size might blur or increase

the size of a spatial representation encoded by a population, but such a representation might be more robust to uncorrelated neural noise because of increased redundancy.

Further complicating matters is the observation that spatial receptive fields have been shown to both increase and decrease in size with attention as a function of where the spatial receptive field is positioned relative to the attended stimulus. Spatial receptive fields tuned near an attended stimulus grow, and spatial receptive fields that fully encompass an attended stimulus shrink^{10,19–23}. These receptive field size changes occur in parallel to changes in the amplitude (gain) of neural responses with attention^{7–17}. Thus, the net impact of all of these changes on the fidelity of population-level spatial representations is unclear, and addressing this issue requires assessing how attention changes the profile of spatial representations encoded by the joint, region-level pattern of activity.

Here we assessed the modulatory role of attention on the spatial information content of putative priority maps by using an encoding model to reconstruct spatial representations of attended and unattended visual stimuli on the basis of multivariate blood oxygen level-dependent functional magnetic resonance imaging (BOLD fMRI) activation patterns within visually responsive regions of the occipital, parietal and frontal cortices. These reconstructions can be considered to reflect region-level spatial representations, and they allowed us to quantitatively track changes in parameters that characterize the topography of spatial maps within each region of interest (ROI). Notably, this technique exploits the full multivariate pattern of BOLD signal across an entire region to evaluate the manner in which spatial representations are modulated by attention rather than comparing multivariate decoding accuracies or considering the univariate response of each voxel in isolation. This approach can be used to examine mechanisms of attentional modulation that cannot be easily

¹Neuroscience Graduate Program, University of California San Diego, La Jolla, California, USA. ²Department of Psychology, University of California San Diego, La Jolla, California, USA. Correspondence should be addressed to T.C.S. (tsprague@ucsd.edu) or J.T.S. (jserences@ucsd.edu).

Received 19 July; accepted 9 October; published online 10 November 2013; doi:10.1038/nn.3574

characterized by measuring changes in either the univariate mean BOLD signal or the decoding accuracy^{24–34} (Fig. 1).

Our results reveal that spatial attention increases the amplitude of region-level stimulus representations within putative priority maps carried by areas of the occipital, parietal and frontal cortices. However, we found little evidence that attention changes the size of stimulus representations in region-level priority maps, even though we observed increases in spatial filter size at the single-voxel level. In addition, the reconstructed spatial representations that are based on activation patterns in later regions of the occipital, parietal and frontal cortices showed larger attentional modulation than those from early areas, which is consistent with the hypothesis that the representations in later regions increasingly transition to more selectively represent relevant stimuli^{4,5}. These changes in the gain of spatial representations should theoretically increase the efficiency with which information about relevant objects in the visual field can be processed and subsequently used to guide perceptual decisions and motor plans¹⁸.

RESULTS

Manipulating attentional demands

To evaluate how task demands influence the topography of spatial representations within different areas of the visual system, we designed a BOLD fMRI experiment that required participants to perform one of three tasks using an identical stimulus display (Fig. 2a). In each trial, participants ($n = 8$) maintained fixation at the center of the screen (Online Methods and Supplementary Fig. 1) while a full-contrast flickering checkerboard was presented in 1 of 36 spatial locations that sampled 6 discrete eccentricities (Fig. 2b). Participants reported either a faint contrast change at the fixation point (the attend fixation condition) or a faint contrast change of the flickering checkerboard stimulus (the attend stimulus condition) or performed a spatial working memory task in which they compared the location of a probe stimulus, T2, with the remembered location of a target stimulus, T1, presented within the radius of the flickering checkerboard (the spatial working memory condition; Fig. 2c). We included the spatial working memory task as an alternate means of inducing focused and sustained spatial attention around the stimulus position³⁵.

On average, performance in the attend fixation task was slightly, although nonsignificantly, higher than that in the attend stimulus or spatial working memory task (Fig. 2d; main effect of condition,

$F(2,14) = 0.951$, $P = 0.41$; attend fixation, $87.37 \pm 6.46\%$ accuracy (mean \pm s.e.m.); attend stimulus, $81.00 \pm 6.67\%$; spatial working memory, $80.00 \pm 2.09\%$). However, we observed a different pattern of response errors across the three task demands: accuracy in the attend fixation condition was lowest in trials in which the flickering checkerboard stimulus was presented near the fixation, whereas accuracy dropped off with increasing stimulus eccentricity in the attend stimulus and spatial working memory tasks (Fig. 2d; condition \times eccentricity interaction, $F(10,70) = 7.235$, $P < 0.0001$).

Reconstructed spatial representations of visual stimuli

To compare spatial representations carried within different brain regions as a function of task demands, we first functionally identified seven ROIs in each hemisphere of each participant using independent localizer techniques (Online Methods and Supplementary Table 1).

Next we used an encoding model^{34,36–38} to reconstruct a spatial representation of the stimulus that was presented during each trial using activation patterns from each ROI (Fig. 3 and Supplementary Figs. 2 and 3). This method results in a spatial representation of the entire visual field measured during each trial that is constrained by activation across all voxels within each ROI. As a result, we obtained average spatial representations for each stimulus position for each ROI for each task condition that accurately reflected the stimulus viewed by the observer (Fig. 4a and Supplementary Fig. 4). This method linearly maps high-dimensional voxel space to a lower-dimensional information space that corresponds to visual field coordinates (Online Methods).

As a point of terminological clarification, we emphasize that we report estimates of the spatial representation of a stimulus display on the basis of the distributed activation pattern across all voxels within a ROI. Throughout the Results section, we therefore refer to our actual measurements as reconstructed spatial representations. However, in the Discussion, we interpret these measurements in the context of putative attentional priority maps that are thought to have a key role in shaping perception and decision making^{1–6}.

Reconstructed spatial representations based on activation patterns in each ROI exhibited several qualitative differences as a function of stimulus eccentricity, task demands and ROI (which we quantify more formally below). First, representations were very precise in the primary visual cortex (V1) (Fig. 4a) and became successively coarser and more diffuse in areas of the extrastriate, parietal and frontal cortices (Fig. 4b).

Figure 1 The effects of spatial attention on region-level priority maps. Spatial attention might act through one of several mechanisms to change the spatial representation of a stimulus within a putative priority map. (a) The hypothetical spatial representation carried across an entire region in response to an unattended circular stimulus. (b) Under one hypothetical scenario, attention might enhance the spatial representation of the same stimulus by amplifying the gain of the spatial representation (i.e., multiplying the representation by a constant greater than 1). (c) Alternatively, attention might act through a combination of multiple mechanisms such as increasing the gain, decreasing the size and increasing the baseline activity of the entire region (i.e., adding a constant to the response across all areas of the priority map). (d) Cross-sections of a–c. This is not meant as an exhaustive description of different attentional modulations. (e) The different types of attentional modulation can give rise to identical responses when the mean BOLD response is measured across the entire expanse of a priority map. Simple Cartesian representations, such as those in a–c, may be visualized in early visual areas where retinotopy is well defined at the spatial resolution of the BOLD response. However, later areas might still encode precise spatial representations of a stimulus even when clear retinotopic organization is not evident, so using alternative methods for reconstructing stimulus representations, such as the approach described in Figure 3, is necessary to evaluate the fidelity of information encoded in putative attentional priority maps. Unattn, unattended; attn, attended.

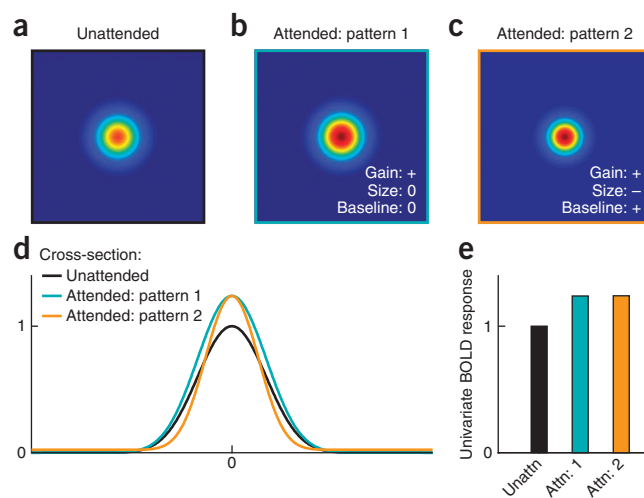
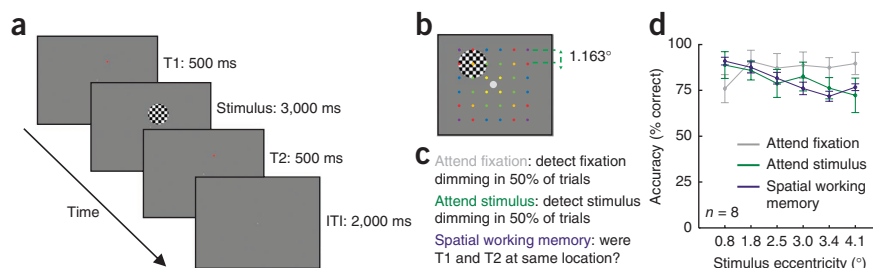


Figure 2 Task design and behavioral results.

(a) Each trial consisted of a 500-ms target stimulus (T1), a 3000-ms flickering checkerboard (6 Hz, full contrast, 2.33° (or 2.326°) diameter) and a 500-ms probe stimulus (T2). T1 and T2 were at the same location in 50% of trials and were slightly offset in the remaining 50% of trials. During the stimulus presentation period, the stimulus dimmed briefly in 50% of trials, and the fixation point dimmed in 50% of trials (each independently randomly chosen).

Participants maintained fixation throughout the experiment, and eye position measured during scanning did not vary as a function of either task demands or stimulus position (**Supplementary Fig. 1**). ITI, intertrial interval. (b) In each trial, a single checkerboard stimulus appeared at 1 of 36 overlapping spatial locations with a slight spatial offset between runs (Online Methods). Each spatial location was sampled once per run. This six-by-six grid of stimulus locations probes six unique eccentricities, as indicated by the color code of the dots (which is not present in the actual stimulus display). (c) In alternating blocks of trials, participants detected either a dimming of the fixation point (attend fixation) or a dimming of the checkerboard stimulus (attend stimulus) or they indicated if the spatial position of T1 and T2 matched (spatial working memory). Notably, all tasks used a physically identical stimulus display—only the task demands varied. Each participant completed between four and six scanning runs of each of the three tasks. (d) For the attend fixation task, performance was better when the stimulus was presented at peripheral locations. In contrast, performance declined with increasing stimulus eccentricity in the attend stimulus and spatial working memory conditions. Error bars, s.e.m.



Similarly, representations of more eccentric stimuli were more diffuse compared to those of more foveal stimuli (for example, when comparing eccentric to foveal representations within each ROI). We also observed higher-fidelity representations of the upper visual field when using only voxels from the ventral aspects of V2 and V3 and higher-fidelity representations of the lower visual field when using only voxels from the dorsal aspects of these regions (**Supplementary Fig. 5a**). These observations, which are consistent with known receptive field properties in nonhuman primates, confirm that our encoding-model method recovered known properties of these visual subregions and these reconstructions were not merely the result of fitting idiosyncratic aspects of our particular data set (i.e., overfitting noise). We further demonstrated this point by using the model to reconstruct representations of completely new stimuli (**Supplementary Fig. 5b**).

Second, the profile of reconstructed spatial representations within many regions also varied with task demands, which is consistent with the notion that these spatial representations reflect spatial maps of attentional priority. Notably, especially in human visual area V4 (hV4),

the human middle temporal cortex (hMT+), the intraparietal sulcus (IPS) and the superior precentral sulcus (sPCS), the magnitude of the spatial representations increased when the participant was either attending to the flickering checkerboard stimulus or performing the spatial working memory task compared to when they were performing a task at fixation.

Size of spatial representations across eccentricities and ROIs

Before formally evaluating the effects of attention on the profile of spatial representations, we first sought to quantify changes in the size of these representations due to stimulus eccentricity and ROI for comparison with known properties of the primate visual system. To this end, we fit a smooth surface to the spatial representations associated with each of the three task conditions separately for each of the 36 possible stimulus locations in each ROI (Online Methods and **Supplementary Fig. 2**). These fits generated an estimate of the amplitude, baseline offset and size of the represented stimulus within each reconstructed spatial representation. We averaged the

Figure 3 The encoding model that was used to reconstruct spatial representations of visual stimuli. Spatial representations of stimuli in each of the 36 possible positions were estimated separately for each ROI.

(a) Training the encoding model. Shown is a set of linear spatial filters that forms the basis set, or information channels, that we used to estimate the spatial selectivity of the BOLD responses in each voxel (Online Methods and **Supplementary Figs. 2 and 3**). The shape of these filters determines how each information channel should respond in each trial given the position of the stimulus that was presented (thus forming a set of regressors or predicted channel responses). Then we constructed a design matrix by concatenating the regressors generated for each trial. This design matrix, in combination with the measured BOLD signal amplitude in each trial, was then used to estimate a weight for each channel in each voxel using a standard general linear model (GLM). (b) Estimating channel responses. Given the known spatial selectivity (or weight) profile of each voxel as computed in a, we then used the pattern of responses across all voxels in each trial in the test set to estimate the magnitude of the response in each of the 36 information channels in that trial. This estimate of the channel responses is thus constrained by the multivariate pattern of responses across all voxels in each trial in the test set and results in a mapping from voxel space (hundreds of dimensions) onto a lower-dimensional channel space (36 dimensions; Online Methods). We then produced a smooth reconstructed spatial representation for every trial by summing the responses of all 36 filters after weighting them by the respective channel responses in each trial. An example of a spatial representation computed from a single trial using data from V1 when the stimulus was presented at the location depicted in a is shown on the lower right.

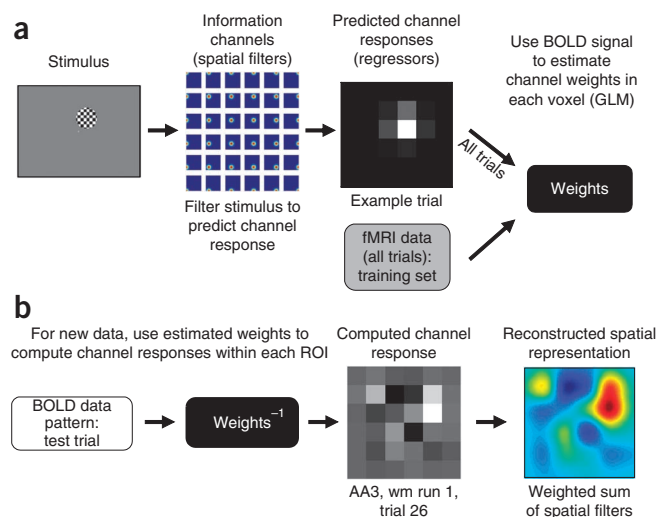
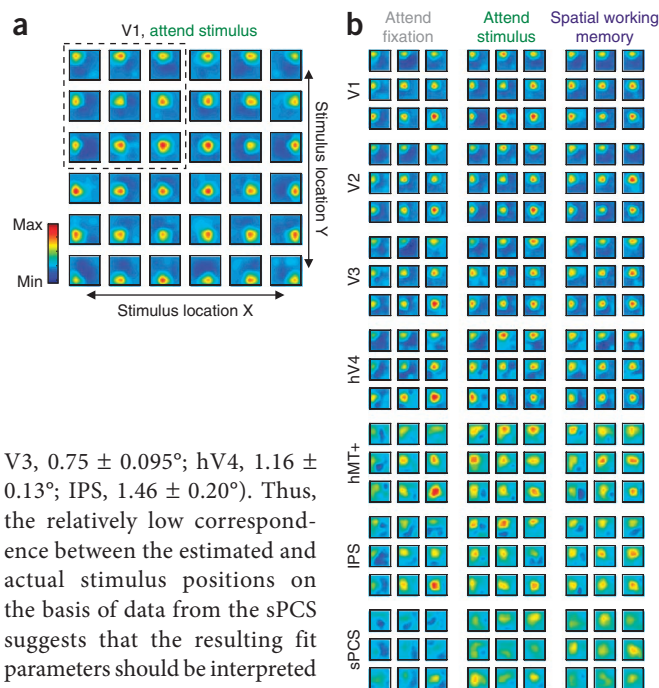


Figure 4 Task demands modulate spatial representations.

(a) Reconstructed spatial representations of each of 36 flickering checkerboard stimuli presented in a six-by-six grid. All 36 stimulus locations are shown, with each location's representation averaged across participants ($n = 8$) using data from bilateral V1 during attend stimulus runs. One participant was not included in this analysis (AG3; **Supplementary Fig. 4**). Each small image represents the reconstructed spatial representation of the entire visual field, and the position of the image corresponds to the location of the presented stimulus. (b) A subset of representations (corresponding to the upper left quadrant of the visual field, represented by the dashed box in a) for each ROI and each task condition. The results were similar for the other quadrants (data not shown; **Fig. 5** shows the aggregate quantification of all reconstructions). All reconstructions in a and b are shown on the same color scale.



fit parameters obtained from each ROI across stimulus locations that were at equivalent eccentricities and then across participants (yielding six sets of fit parameters, with one set for each of the six possible stimulus eccentricities; **Fig. 2b**). We then used these fit parameters to make inferences about how the magnitudes and shapes of the spatial representations of stimuli from each ROI varied across stimulus positions.

First we quantified the accuracy of fits by computing the Euclidean distance between the centroid of the fit function and the actual location of the stimulus across all eccentricities and task conditions. The estimated centroids were generally accurate and closely tracked changes in stimulus location (**Fig. 4a**). However, the distances between the fit centroids and the actual stimulus positions in sPCS were nearly double those of the next least-accurate region, hMT+ (sPCS, $3.01 \pm 0.077^\circ$ (mean \pm s.e.m.); hMT+, $1.68 \pm 0.17^\circ$). The error distances in all other areas were relatively small (V1, $0.67 \pm 0.084^\circ$; V2, $0.77 \pm 0.12^\circ$;

V3, $0.75 \pm 0.095^\circ$; hV4, $1.16 \pm 0.13^\circ$; IPS, $1.46 \pm 0.20^\circ$). Thus, the relatively low correspondence between the estimated and actual stimulus positions on the basis of data from the sPCS suggests that the resulting fit parameters should be interpreted with caution (addressed further in the Discussion).

In early visual ROIs of V1, V2, V3 and hV4, the size of the reconstructed spatial representations increased with increasing eccentricity regardless of task condition (**Fig. 5** and Online Methods; main effect of eccentricity, two-way analysis of variance (ANOVA) within each ROI, all $P < 0.0004$; unless otherwise specified, all statistical tests on

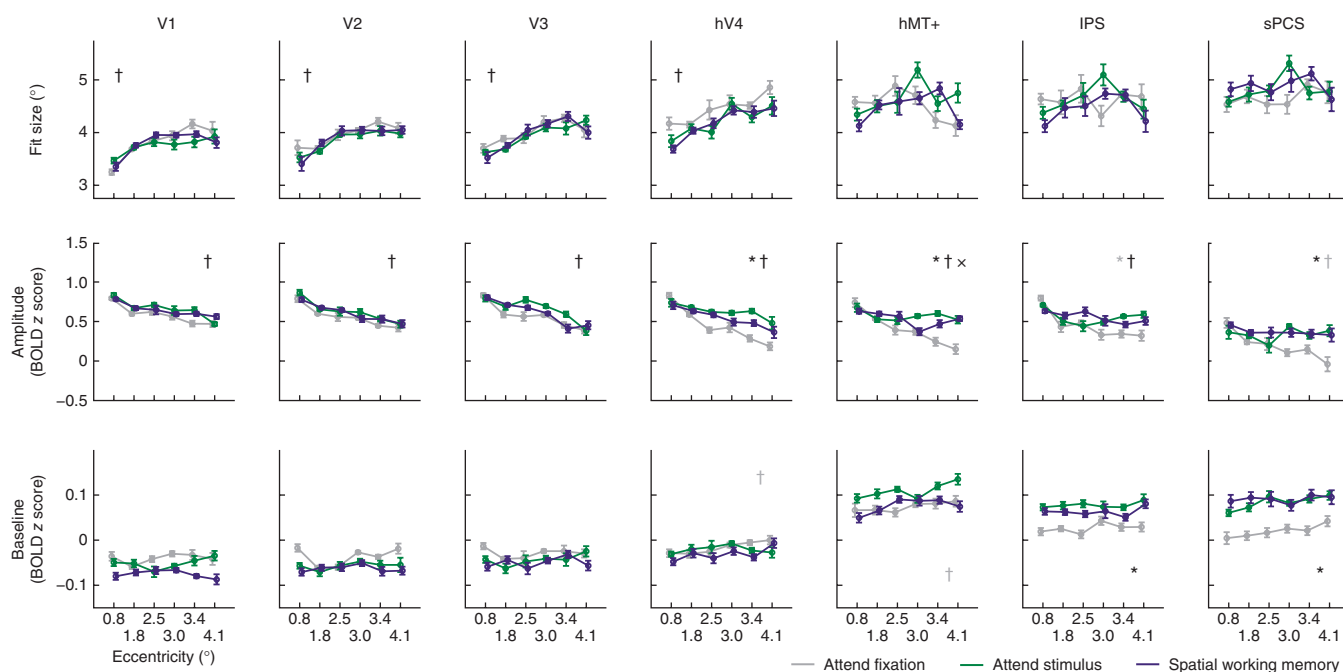
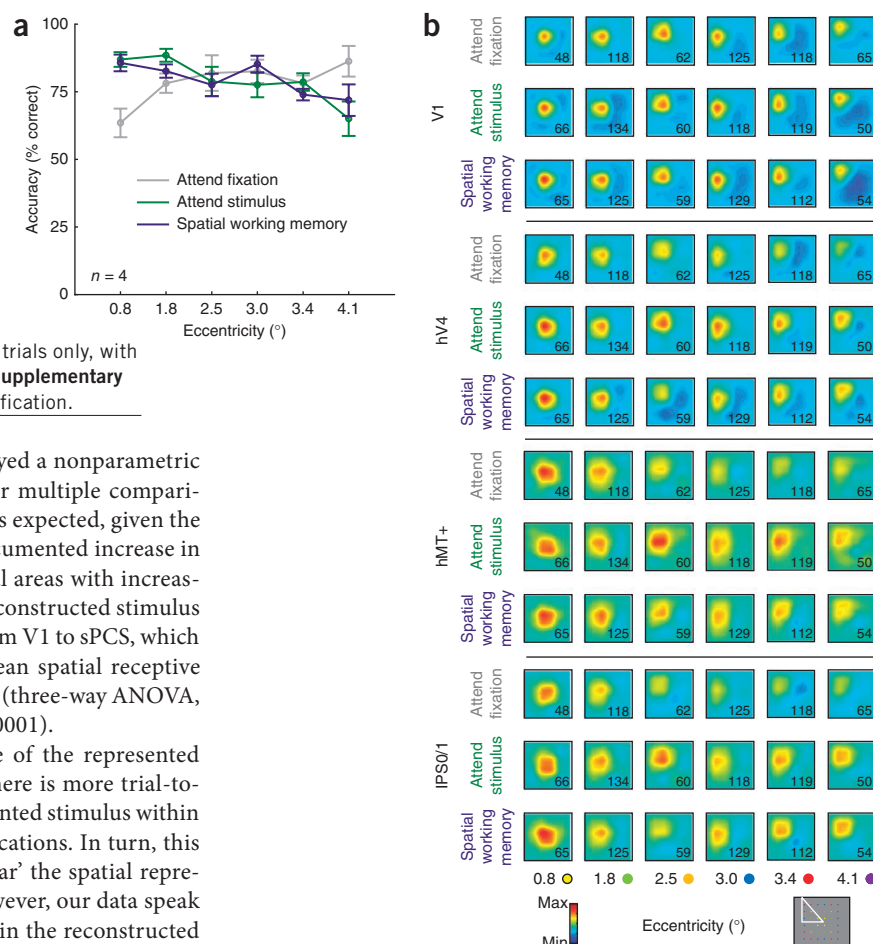


Figure 5 Fit parameters to reconstructed spatial representations averaged across like eccentricities. For each participant, we fit a smooth two-dimensional surface (Online Methods) to the average reconstructed stimulus representation in all 36 locations separately for each task condition and ROI. We allowed the amplitude, baseline, size and center (x, y coordinate) of the fit basis function to vary freely during the fitting. Fit parameters were averaged within each participant across like eccentricities and then averaged across participants. The size of the best-fitting surface varied systematically with stimulus eccentricity and ROI but did not vary as a function of task condition. In contrast, the amplitude of the best fitting surface increased with attention in hV4, hMT+ and sPCS (with a marginal effect in IPS). Shown are the main effect of task condition (*), eccentricity (†) and interaction between task and eccentricity (x) at the $P < 0.05$ level corrected for multiple comparisons (Online Methods). Gray symbols indicate trends at the $P < 0.025$ level uncorrected for multiple comparisons. Error bars, within-participant s.e.m.

Figure 6 Results are consistent when task difficulty is matched. **(a)** Performance of four participants who were rescanned while carefully matching task difficulty across all three experimental conditions. As in **Figure 2d**, the participants' performance is better in the attend fixation task when the checkerboard is presented in the periphery, and performance in the attend stimulus and spatial working memory tasks is better when the stimulus is presented near the fovea. **(b)** A subset of illustrative reconstructed stimulus representations from V1, hV4, hMT+ and IPS0 and IPS1 (IPS0/1) averaged across like eccentricities (correct trials only, with the number of averaged trials indicated as insets). **Supplementary Figure 7** includes details about IPS subregion identification.



fit parameters to spatial representations employed a nonparametric permutation procedure and were corrected for multiple comparisons). This increase in size with eccentricity was expected, given the use of a constant stimulus size and the well-documented increase in the size of spatial receptive fields in early visual areas with increasing eccentricity³⁹. In addition, the size of the reconstructed stimulus representations also increased systematically from V1 to sPCS, which is consistent with the known expansion of mean spatial receptive field sizes in the parietal and frontal cortices^{40,41} (three-way ANOVA, significant main effect of ROI on fit size, $P < 0.0001$).

One alternative explanation is that the size of the represented stimulus increases with eccentricity because there is more trial-to-trial variability in the center point of the represented stimulus within reconstructions at more peripheral stimulus locations. In turn, this increase in trial-to-trial variability would 'smear' the spatial representations, leading to larger size estimates. However, our data speak against this possibility, as increased variability in the reconstructed stimulus locations would also result in lower estimated amplitudes, so increases in fit size and decreases in fit amplitude across conditions would always be yoked, and correlating the change in amplitude and the change in size within each eccentricity across each condition pair would reveal a negative correlation (for example, if the size of the spatial representation measured at a given eccentricity increased with attention, then the amplitude would decrease). No combinations of condition pair, eccentricity and ROI revealed a significant correlation between change in amplitude and change in size (all $P > 0.05$, corrected using the false discovery rate (FDR); Online Methods). Furthermore, in a follow-up analysis, we computed the population receptive field (pRF) for each voxel⁴², which revealed that voxels tuned to more eccentric visual field positions have a larger pRF size (**Supplementary Table 2, Supplementary Results** and Online Methods). This combination of analyses supports the conclusion that increases in fit size with increases in stimulus eccentricity are not due solely to increased variability in reconstructed spatial representations.

Effects of attention on spatial representations

Despite being sensitive to expected changes in representation size on the basis of anatomical properties of the visual system, task demands exerted a negligible influence on the size of the reconstructed spatial representations, with no areas showing a significant effect (hV4 was closest at $P = 0.033$, but this did not survive correction for multiple comparisons, and P values in all other regions were >0.147).

In contrast, the fit amplitudes in hV4, hMT+, IPS and sPCS were significantly modulated by task condition, with a higher amplitude in the attention and working memory conditions than in the fixation condition (**Fig. 5**; three-way ANOVA, main effect of task condition,

$P = 0.0003$). For example, in hV4, the amplitude of the best fitting surface to the spatial representations of attended stimuli was higher during the attend stimulus and spatial working memory conditions as compared to the attend fixation condition (two-way ANOVA, main effect of task condition, $P < 0.0001$). We observed similar effects in hMT+ (two-way ANOVA, $P = 0.0007$) and sPCS (two-way ANOVA, $P = 0.0007$). A similar pattern was evident in IPS as well, but it did not survive correction for multiple comparisons (two-way ANOVA, uncorrected $P = 0.011$). Within individual ROIs, there was a significant interaction between task condition and eccentricity in hMT+ ($P = 0.0003$), with larger increases in amplitude observed for more eccentric stimuli. It is notable that this increase in the amplitude of spatial representations with attention corresponds to a focal gain modulation that is restricted to the portion of visual space in the immediate neighborhood of the attended stimulus. Changes in fit amplitude do not result from a uniform, region-wide increase in BOLD signal that equally influences the response across an entire ROI; such a general and widespread modulation would be accounted for by an increase in the baseline fit parameter (**Supplementary Fig. 2**). In addition, the impact of task condition on the amplitude of reconstructed spatial representations was more pronounced in later visual areas (hV4, hMT+, IPS and sPCS) compared to earlier areas (V1, V2 and V3) (three-way interaction between ROI, condition and eccentricity, $P = 0.043$).

In addition to an increase in the fit amplitude of the reconstructed spatial representations, IPS and sPCS also exhibited a spatially global increase in baseline response levels across the entire measured spatial representation in the attend stimulus and spatial working memory conditions compared to the attend fixation condition (**Fig. 5** and

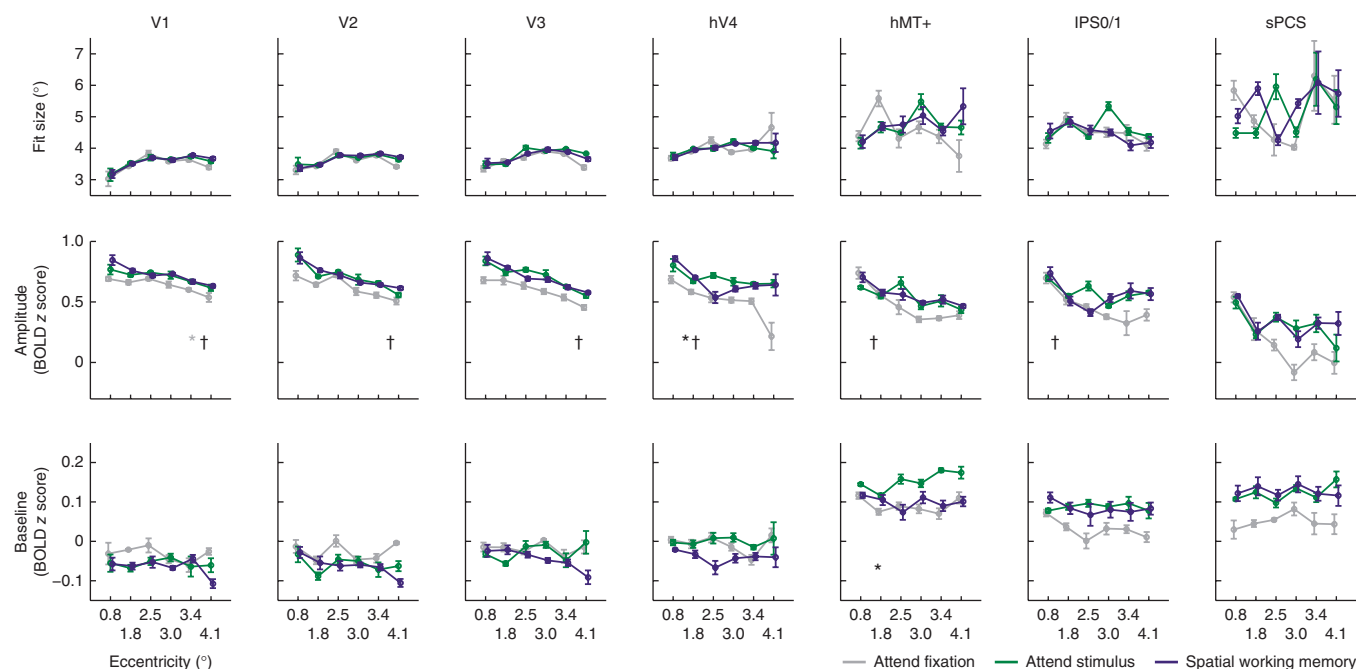


Figure 7 Fit parameters to spatial representations after controlling for task difficulty. As in **Figure 5**, a surface was fit to the averaged, co-registered spatial representations for each participant. However, in this case, task difficulty was carefully matched between conditions, and representations were based solely on trials in which the participant made a correct behavioral response (**Fig. 6b**). The results are similar to those reported in **Figure 5**: attention acts to increase the fit amplitude of spatial representations in hV4 but does not act to decrease size. In hMT+, attention also acted in a nonlocalized manner to increase the baseline parameter. The statistics and symbols are as in **Figure 5**. Error bars, within-participant s.e.m.

Supplementary Fig. 6; two-way ANOVA, main effect of condition, IPS, $P = 0.0014$; sPCS, $P = 0.0012$). The spatially nonselective increases may reflect the fact that spatial receptive fields in these regions are often large enough to encompass the entire stimulus display^{40,41}, so all stimuli might drive some increase in the response irrespective of spatial position.

Controlling difficulty across task conditions

Slight differences in task difficulty in the first experiment (**Fig. 2d**) might have contributed to the observed changes in spatial representations. To address this possibility, we ran four participants from the original cohort in a second experimental session while carefully equating behavioral performance across all three tasks (**Fig. 6a**). Overall accuracy during this second session did not differ significantly across the three conditions, although we observed a similar interaction between task condition and stimulus eccentricity (**Fig. 6a**; two-way repeated-measures ANOVA, main effect of condition, $F(2,6) = 0.043$, $P = 0.96$; condition \times eccentricity interaction, $F(10,30) = 3.28$, $P = 0.005$; attend fixation, $78.8 \pm 2.80\%$ (mean \pm s.e.m.); attend stimulus, $80.0 \pm 2.60\%$; spatial working memory, $79.8 \pm 1.76\%$). In addition, we also identified IPS visual field maps 0–3 (IPS0–IPS3) using standard procedures so that we could more precisely characterize the effects of attention on stimulus representations in subregions of our larger IPS ROI^{31,32,43,44} (Online Methods and **Supplementary Fig. 7**).

To ensure that behavioral performance was not unduly biasing our results, we reconstructed spatial representations using only correct trials ($\sim 80\%$ of total trials; **Fig. 6a**). All representations were co-registered on the basis of stimulus eccentricity before averaging (the corresponding eccentricity points are shown in **Fig. 2b**). Even though our sample size was smaller ($n = 4$ as compared to $n = 8$), the influence of attention on the topography of the spatial representations was similar to that in our initial observations (**Fig. 6b**). In addition, mapping out retinotopic subregions of the IPS revealed that the functionally

defined IPS ROI (shown in **Fig. 5**) corresponds primarily to IPS0 and IPS1 (**Supplementary Fig. 7a,b**).

When examining best-fit surfaces to the spatial representations from this experiment (we computed the fits using co-registered representations and only correct trials for each participant; **Fig. 7** and Online Methods), we found that attention significantly modulated amplitude across all regions (three-way ANOVA, main effect of task condition, $P = 0.0162$). When considered in isolation, only hV4 showed a significant change in amplitude with attention after correction for multiple comparisons (two-way repeated-measures ANOVA, $P = 0.0022$). However, we observed similar trends in V1, V2 and V3 (uncorrected $P = 0.0243$, 0.042 and 0.031 , respectively). We found no significant main effect of task condition on the size of the representations (all $P > 0.135$, with the minimum P found for hMT+), and the overall baseline levels significantly increased as a function of task condition in hMT+ only ($P = 0.00197$). Across all ROIs, there was a main effect of eccentricity on fit size (three-way ANOVA, $P = 0.0016$) but no main effect of task condition on fit size (three-way ANOVA, $P = 0.423$).

pRFs expand with attention

For these same four participants, we computed the pRF⁴² for each voxel in V1, hV4, hMT+ and IPS0 using data from the behaviorally controlled replication experiment. We computed pRFs by first using the initial step of our encoding-model estimation procedure (**Fig. 3a**) to determine the response of each voxel to each position in the visual field (**Supplementary Figs. 8 and 9** and Online Methods). We then fit each voxel's response profile with the same surface that we used to characterize the spatial representations. By comparing pRFs computed using data from each condition independently, we found that a majority of the pRFs in hV4, hMT+ and IPS0 increased in size during either the attend stimulus or spatial working memory condition as

compared to the attend fixation condition. In contrast, pRF size in V1 was not significantly modulated by attention (**Supplementary Fig. 9** and **Supplementary Results**).

To reconcile the results that voxel-level pRFs expanded with attention yet region-level spatial representations remained at a constant size, we simulated data using estimated pRF parameters from hV4 (a region for which spatial representations increase in amplitude and pRFs increase in size; Online Methods) under different pRF modulation conditions. In the first condition, we generated simulated data using pRFs with sizes centered around two mean values, which resulted in a pRF scaling across all simulated voxels (the average size across voxels increased, but some voxels decreased in size and others increased). Under these conditions, spatial representations increased in size (**Supplementary Fig. 10a,b**). In a second pRF modulation scenario, we used the fit pRF values from one participant's hV4 ROI (**Supplementary Fig. 8**) to simulate data. In this case, spatial representations remained the same size but increased in amplitude, which is consistent with our observations using real data (**Figs. 5** and **7** and **Supplementary Fig. 10c,d**; this conclusion was also supported when we used pRF data from the other three observers to seed the simulation). Thus, the pattern of pRF modulations across all voxels enhances the amplitude of spatial representations while preserving their size.

DISCUSSION

Spatial attention has previously been shown to alter the gain of single-unit responses that are associated with relevant visual features such as orientation^{7–9,12,13,16,17} and motion direction^{11,14,15}, as well as modulate the size of spatial receptive fields^{10,19–23}. Here we show that these local modulations operate jointly to increase the overall amplitude of the region-level spatial representation of an attended stimulus without changing its represented size. Furthermore, these amplitude modulations were especially apparent in later areas of the visual system such as hV4, hMT+ and IPS, which is consistent with predictions made by computational theories of attentional priority maps^{4,5}.

We were able to reconstruct robust spatial representations across a range of eccentricities and for all three task conditions in all measured ROIs. Notably, even though we used an identical reconstruction procedure in all areas, the size of the reconstructed spatial representations increased from early to later visual areas (**Fig. 5**). Single-unit receptive field sizes across cortical regions are thought to increase in a similar manner^{39–41,45,46}. In addition, representations of stimuli presented at higher eccentricities were larger than representations of stimuli presented near the fovea, which also corresponds to known changes in receptive field size with eccentricity^{39,42}. Furthermore, simulating data under conditions in which we uniformly scaled the mean size of voxel-level pRFs revealed that such changes are detectable using our analysis method (**Supplementary Fig. 10a,b**). Thus, this technique is sensitive to detecting changes in the size of spatial representations of stimuli that are driven by known neural constraints such as relative differences in receptive field size across cortical ROIs and eccentricities, even though these factors are not built in to the spatial encoding model. Together these empirical and modeling results suggest that at the level of region-wide priority maps, the representation of a stimulus does not expand or contract under the attentional conditions tested here, and they underscore the importance of incorporating response changes across all encoding units when evaluating attentional modulations.

The quantification method we implemented for measuring changes in spatial representations across tasks, eccentricities and ROIs involved fitting a surface that was defined by several parameters: center location, amplitude, baseline offset and size (**Supplementary Fig. 2**).

Changes in activation that carry no information about stimulus location (such as changes in general arousal or responsiveness to stimuli presented in all locations because of large receptive fields) will influence the baseline parameter, as such changes reflect increased or decreased signal across an entire region. In contrast, a change in the spatial representation that changes the representation of a visual stimulus would result in a change in the amplitude or size parameter (or both). Here we demonstrated that attention operates primarily by selectively increasing the amplitude of stimulus representations in several putative priority maps (**Figs. 5** and **7**) rather than increasing the overall BOLD signal more generally across entire regions.

Notably, spatial reconstructions based on activation patterns from sPCS were relatively inaccurate compared to other ROIs, and this ROI primarily exhibited increases in the fit baseline parameter (**Fig. 5**). This region, which may be a human homolog of the functionally defined macaque frontal eye field^{47,48} (FEF), might have showed degraded spatial selectivity in the present study because of the relatively large size of spatial receptive fields observed in many FEF neurons (typically $\geq 20^\circ$ diameter⁴¹) and the small area subtended by our stimulus display (9.31° across horizontally). Consistent with this possibility, previous reports of retinotopic organization in the human frontal cortex used stimuli that were presented at higher eccentricities in order to resolve spatial maps ($\geq 10^\circ$ (ref. 49) to 25° (ref. 45)).

Attentional priority maps

The extensive literature on spatial salience or priority maps^{1–6} postulates the existence of one or several maps of visual space, each carrying information about behaviorally relevant objects within the visual scene. Furthermore, priority maps in early visual areas (for example, V1) are thought to encode primarily low-level stimulus features (for example, contrast), whereas priority maps in later regions are thought to increasingly weight behavioral relevance over low-level stimulus attributes⁴. Although many important insights have stemmed from observing single-unit responses as a function of changes in attentional priority (reviewed in ref. 5), these results provide information about how isolated pixels in a priority map change under different task conditions.

A previous fMRI study used multivariate decoding (classification) analyses to identify several frontal and parietal ROIs that exhibit similar activation patterns during covert attention, spatial working memory and saccade generation tasks³². These results provide strong support for the notion that common priority maps support representations of attentional priority across multiple tasks. Here we assessed how the holistic landscape across these priority maps measured using fMRI changed as attention was systematically varied. Our demonstration that spatial representation amplitude was enhanced with attention in later, but not earlier, ROIs supports the hypothesis that priority maps in higher areas are increasingly dominated by attentional factors and suggests that these attentional modulations of priority maps operate by scaling the amplitude of the behaviorally relevant item without changing its represented size.

pRFs

In addition to measuring spatial representations that are carried by the pattern of activation across entire visual regions, we also estimated the voxel-level pRFs⁴² for a subset of participants and ROIs by adding constraints to our encoding-model estimation procedure (**Supplementary Figs. 8** and **9** and Online Methods). This alternative tool has been used previously to evaluate the aggregate spatial receptive field profile across all neural populations within voxels that belong to different visual ROIs⁴².

Changes in voxel-level pRFs can inform how a region dynamically adjusts the spatial sensitivity of its constituent filters in order to modulate its overall spatial priority map. First, we replicated the typical results that voxel-level pRFs tuned for more eccentric visual field positions are larger in size (**Supplementary Table 2**) and that pRFs for later visual regions tend to be larger than pRFs for earlier visual regions (**Supplementary Fig. 9**). Second, results from this complementary analysis revealed that in regions that showed enhanced spatial representation amplitude with attention (hV4, hMT+ and IPS0), pRF size increased (**Supplementary Figs. 8 and 9**), even though the corresponding region-level spatial representations did not increase in size (**Fig. 7**). This may seem like a disconnect, given that the particular pattern of pRF changes across all voxels within a region jointly shapes how the spatial priority map changes with attention. However, there is not necessarily a monotonic mapping between the size of the constituent filters and the size of population-level spatial representations. Indeed, simulations based on the observed pattern of pRF changes with attention give rise to region-level increases in representation amplitude in the absence of changes in representation size, as we observed in our data (**Supplementary Fig. 10**). This finding, together with our primary results concerning region-level spatial representations, provides evidence that attentional modulation of spatial information encoding is a process that strongly benefits from study at the large-scale population level.

Comparison to previous results

At the level of single-unit recordings, attention has been shown to decrease the size of MT spatial receptive fields when an animal is attending to a stimulus that is encompassed by the recorded neuron's receptive field^{19–21} and to increase the size of spatial receptive fields when an animal is attending nearby the recorded neuron's receptive field^{20–22}. In V4, spatial receptive fields appear to shift toward the attended region of space in a subset of neurons¹⁰. With respect to cortical space, these single-unit attentional modulations of spatial receptive fields suggest that unifocal attention may act to increase the cortical surface area that is responsive to a stimulus of constant size. Consistent with this prediction, our measured pRFs for extrastriate regions of hV4, hMT+ and IPS0 increased in size with attention.

In contrast, one previous report suggested that spatial attention instead narrows the activation profile along the cortical surface of the visual cortex in response to a visual stimulus⁵⁰. However, this inference was based on patterns of intertrial correlations between BOLD activation patterns that were associated with dividing attention between four stimuli (one presented in each quadrant). These patterns were suggested to result from a combination of attention-related gain and narrowing of population-level responses⁵⁰, that is, a narrower response along the cortical surface with attention.

We did not observe any significant attention-related changes in the size of the reconstructed spatial representations in either the primary visual cortex or other areas in the extrastriate, parietal or frontal cortices. However, the tasks performed by observers and the analysis techniques implemented were very different between these studies. Most notably, observers in the present study and in previous fMRI^{24–33} and single-unit studies^{10,19–21} were typically required to attend to a single stimulus, whereas population-level activation narrowing was observed when participants simultaneously attended to the precise spatial position of four Gabor stimuli, one presented in each visual quadrant⁵⁰. Furthermore, our observation that pRFs increased in size during the attend stimulus and spatial working memory conditions is compatible with the pattern of spatial receptive field changes in

single units^{10,19–23}, and our data and simulations show that these local changes can result in a region-level representation that changes only in amplitude and not in size (**Supplementary Fig. 10**).

Collectively, it seems probable that the exact task demands (unifocal as compared to multifocal attention) and stimulus properties (single stimulus as compared to multiple stimuli) may have a key role in determining how attention influences the profile of spatial representations. Future work using analysis methods that are sensitive to region-level differences in spatial representations (for example, applying encoding models similar to that described here to data acquired when participants perform different tasks) in conjunction with careful identification of neural receptive field properties across those task-demand conditions (for example, from simultaneous multiunit electrophysiological recordings or *in vivo* two-photon Ca²⁺ imaging in rodents and primates) may provide complementary insights into when and how attention changes the shape and/or amplitude of stimulus representations in spatial priority maps and how those changes are implemented in the neural circuitry.

Notably, although our observations are largely consistent with measured receptive field changes at the single-unit level^{10,19–23}, we cannot make direct inferences that such single-unit changes are in fact occurring. A number of mechanisms, including one in which only the gain of different populations is modulated by attention, could also account for the pattern of results we saw in both our region-level spatial representations (**Figs. 5 and 7**) and our pRF measurements (**Supplementary Figs. 8 and 9**). We note, however, that some neural mechanisms are highly unlikely given our measured spatial representations and pRFs. For example, we would not observe an increase in pRF size if spatial receptive fields of neurons within those voxels were to exclusively narrow with attention. As a result of these interpretational concerns, we restrict the inferences we draw from our results to the role of attention in modulating region-level spatial priority maps measured with fMRI and make no direct claims about spatial information coding at a neural level.

Information content of attentional priority maps

One consequence of an observed increase in the amplitude of reconstructed priority maps is that the mutual information between the stimulus position and the observed BOLD responses should increase (a more complete discussion is provided in ref. 18). This increase can occur, in theory, because mutual information reflects the ratio of signal entropy (the variability in neural responses that is tied systematically to changes in the stimulus) to noise entropy (the variability in neural responses that is not tied to changes in the stimulus). Thus a multiplicative increase in the gain of the neural responses that are associated with an attended stimulus should increase mutual information because it will increase the variability of responses that are associated with an attended stimulus location, which will in turn increase signal entropy. In contrast, a purely additive shift in all neural responses (reflected by an increase in the fit baseline parameter) will not increase the dynamic range of responses that are associated with an attended stimulus location, causing the mutual information to either remain constant (under a constant additive noise model) or even decrease (under a Poisson noise model, in which noise increases with the mean). Previous fMRI work on spatial attention has not attempted to disentangle these two potential sources of increases in the BOLD signal, highlighting the utility of approaches that can support more precise inferences about how task demands influence region-level neural codes^{24–33}.

The information content of a neural code is not necessarily monotonically related to the size of the constituent neural filters¹⁸.

Extremely small (pinpoint) or extremely large (flat) spatial filters each individually carry very little information about the spatial arrangement of stimuli within the visual field. Accordingly, the optimal filter size lies somewhere between these two extremes, and thus it is not straightforward to infer whether a change in filter size results in a more or less optimal neural code (in terms of information encoding capacity). By simultaneously estimating changes in filter size across an entire ROI subtending the entire stimulated visual field, we were able to demonstrate that the synergistic pattern of spatial filter (pRF) modulations with attention jointly constrains the region-level spatial representation to maintain a constant represented stimulus size, despite most voxels exhibiting an increase in pRF size (**Supplementary Figs. 8–10**). Together our results demonstrate the importance of incorporating all available information across entire ROIs when evaluating the modulatory role of attention on the information content of spatial priority maps.

METHODS

Methods and any associated references are available in the [online version of the paper](#).

Note: Any Supplementary Information and Source Data files are available in the online version of the paper.

ACKNOWLEDGMENTS

We thank E. Vul and S. Itthipuripat for assistance with statistical methods and M. Scolari and M. Smith for assistance with parietal cortex mapping protocols. This work was supported by a National Science Foundation Graduate Research Fellowship to T.C.S. and by US National Institutes of Health grant R01 MH-092345 and a James S. McDonnell Scholar Award to J.T.S.

AUTHOR CONTRIBUTIONS

T.C.S. and J.T.S. designed the experiments and analysis method and wrote the manuscript. T.C.S. conducted the experiments and implemented the analyses.

COMPETING FINANCIAL INTERESTS

The authors declare no competing financial interests.

Reprints and permissions information is available online at <http://www.nature.com/reprints/index.html>.

- Koch, C. & Ullman, S. Shifts in selective visual attention: towards the underlying neural circuitry. *Hum. Neurobiol.* **4**, 219–227 (1985).
- Itti, L., Koch, C. & Niebur, E. A model of saliency-based visual attention for rapid scene analysis. *IEEE Trans. Pattern Anal. Mach. Intell.* **20**, 1254–1259 (1998).
- Itti, L. & Koch, C. Computational modelling of visual attention. *Nat. Rev. Neurosci.* **2**, 194–203 (2001).
- Serences, J.T. & Yantis, S. Selective visual attention and perceptual coherence. *Trends Cogn. Sci.* **10**, 38–45 (2006).
- Fecteau, J.H. & Munoz, D.P. Saliency, relevance, and firing: a priority map for target selection. *Trends Cogn. Sci.* **10**, 382–390 (2006).
- Bichot, N.P. & Schall, J.D. Effects of similarity and history on neural mechanisms of visual selection. *Nat. Neurosci.* **2**, 549–554 (1999).
- Luck, S.J., Chelazzi, L., Hillyard, S.A. & Desimone, R. Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. *J. Neurophysiol.* **77**, 24–42 (1997).
- Reynolds, J.H., Chelazzi, L. & Desimone, R. Competitive mechanisms subserve attention in macaque areas V2 and V4. *J. Neurosci.* **19**, 1736–1753 (1999).
- McAdams, C.J. & Maunsell, J.H.R. Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *J. Neurosci.* **19**, 431–441 (1999).
- Connor, C.E., Preddie, D.C., Gallant, J.L. & Van Essen, D.C. Spatial attention effects in macaque area V4. *J. Neurosci.* **17**, 3201–3214 (1997).
- Treue, S. & Maunsell, J.H.R. Attentional modulation of visual motion processing in cortical areas MT and MST. *Nature* **382**, 539–541 (1996).
- Reynolds, J.H., Pasternak, T. & Desimone, R. Attention increases sensitivity of V4 neurons. *Neuron* **26**, 703–714 (2000).
- McAdams, C.J. & Maunsell, J.H.R. Attention to both space and feature modulates neuronal responses in macaque area V4. *J. Neurophysiol.* **83**, 1751–1755 (2000).
- Treue, S. & Maunsell, J.H.R. Effects of attention on the processing of motion in macaque middle temporal and medial superior temporal visual cortical areas. *J. Neurosci.* **19**, 7591–7602 (1999).
- Seidemann, E. & Newsome, W.T. Effect of spatial attention on the responses of area MT neurons. *J. Neurophysiol.* **81**, 1783–1794 (1999).
- Motter, B.C. Focal attention produces spatially selective processing in visual cortical areas V1, V2, and V4 in the presence of competing stimuli. *J. Neurophysiol.* **70**, 909–919 (1993).
- Moran, J. & Desimone, R. Selective attention gates visual processing in the extrastriate cortex. *Science* **229**, 782–784 (1985).
- Saproo, S. & Serences, J.T. Spatial attention improves the quality of population codes in human visual cortex. *J. Neurophysiol.* **104**, 885–895 (2010).
- Womelsdorf, T., Anton-Erxleben, K., Pieper, F. & Treue, S. Dynamic shifts of visual receptive fields in cortical area MT by spatial attention. *Nat. Neurosci.* **9**, 1156–1160 (2006).
- Womelsdorf, T., Anton-Erxleben, K. & Treue, S. Receptive field shift and shrinkage in macaque middle temporal area through attentional gain modulation. *J. Neurosci.* **28**, 8934–8944 (2008).
- Anton-Erxleben, K., Stephan, V.M. & Treue, S. Attention reshapes center-surround receptive field structure in macaque cortical area MT. *Cereb. Cortex* **19**, 2466–2478 (2009).
- Niebergall, R., Khayat, P.S., Treue, S. & Martinez-Trujillo, J.C. Expansion of MT neurons excitatory receptive fields during covert attentive tracking. *J. Neurosci.* **31**, 15499–15510 (2011).
- Anton-Erxleben, K. & Carrasco, M. Attentional enhancement of spatial resolution: linking behavioural and neurophysiological evidence. *Nat. Rev. Neurosci.* **14**, 188–200 (2013).
- Liu, T., Pestilli, F. & Carrasco, M. Transient attention enhances perceptual performance and fMRI response in human visual cortex. *Neuron* **45**, 469–477 (2005).
- Gandhi, S.P., Heeger, D.J. & Boynton, G.M. Spatial attention affects brain activity in human primary visual cortex. *Proc. Natl. Acad. Sci. USA* **96**, 3314–3319 (1999).
- Kastner, S., Pinsk, M.A., De Weerd, P., Desimone, R. & Ungerleider, L.G. Increased activity in human visual cortex during directed attention in the absence of visual stimulation. *Neuron* **22**, 751–761 (1999).
- Brefczynski, J.A. & DeYoe, E.A. A physiological correlate of the “spotlight” of visual attention. *Nat. Neurosci.* **2**, 370–374 (1999).
- Silver, M.A., Ress, D. & Heeger, D.J. Neural correlates of sustained spatial attention in human early visual cortex. *J. Neurophysiol.* **97**, 229–237 (2007).
- Tootell, R.B. *et al.* The retinotopy of visual spatial attention. *Neuron* **21**, 1409–1422 (1998).
- Murray, S.O. The effects of spatial attention in early human visual cortex are stimulus independent. *J. Vis.* **8**, 2.1–2.11 (2008).
- Silver, M.A., Ress, D. & Heeger, D.J. Topographic maps of visual spatial attention in human parietal cortex. *J. Neurophysiol.* **94**, 1358–1371 (2005).
- Jerde, T.A., Merriam, E.P., Riggall, A.C., Hedges, J.H. & Curtis, C.E. Prioritized maps of space in human frontoparietal cortex. *J. Neurosci.* **32**, 17382–17390 (2012).
- Jehee, J.F.M., Brady, D.K. & Tong, F. Attention improves encoding of task-relevant features in the human visual cortex. *J. Neurosci.* **31**, 8210–8219 (2011).
- Serences, J.T. & Saproo, S. Computational advances towards linking BOLD and behavior. *Neuropsychologia* **50**, 435–446 (2012).
- Awh, E. & Jonides, J. Overlapping mechanisms of attention and spatial working memory. *Trends Cogn. Sci.* **5**, 119–126 (2001).
- Brouwer, G.J. & Heeger, D.J. Decoding and reconstructing color from responses in human visual cortex. *J. Neurosci.* **29**, 13992–14003 (2009).
- Naselaris, T., Kay, K., Nishimoto, S. & Gallant, J. Encoding and decoding in fMRI. *Neuroimage* **56**, 400–410 (2011).
- Scolari, M., Byers, A. & Serences, J.T. Optimal deployment of attentional gain during fine discriminations. *J. Neurosci.* **32**, 7723–7733 (2012).
- Gattass, R. *et al.* Cortical visual areas in monkeys: location, topography, connections, columns, plasticity and cortical dynamics. *Phil. Trans. R. Soc. Lond. B* **360**, 709–731 (2005).
- Ben Hamed, S., Duhamel, J.R., Bremmer, F. & Graf, W. Visual receptive field modulation in the lateral intraparietal area during attentive fixation and free gaze. *Cereb. Cortex* **12**, 234–245 (2002).
- Mohler, C.W., Goldberg, M.E. & Wurtz, R.H. Visual receptive fields of frontal eye field neurons. *Brain Res.* **61**, 385–389 (1973).
- Dumoulin, S.O. & Wandell, B.A. Population receptive field estimates in human visual cortex. *Neuroimage* **39**, 647–660 (2008).
- Sereno, M.I., Pitzalis, S. & Martinez, A. Mapping of contralateral space in retinotopic coordinates by a parietal cortical area in humans. *Science* **294**, 1350–1354 (2001).
- Swisher, J.D., Halko, M.A., Merabet, L.B., McMains, S.A. & Somers, D.C. Visual topography of human intraparietal sulcus. *J. Neurosci.* **27**, 5326–5337 (2007).
- Saygin, A.P. & Sereno, M.I. Retinotopy and attention in human occipital, temporal, parietal, and frontal cortex. *Cereb. Cortex* **18**, 2158–2168 (2008).
- Kastner, S. *et al.* Modulation of sensory suppression: implications for receptive field sizes in the human visual cortex. *J. Neurophysiol.* **86**, 1398–1411 (2001).
- Srimal, R. & Curtis, C.E. Persistent neural activity during the maintenance of spatial position in working memory. *Neuroimage* **39**, 455–468 (2008).
- Paus, T. Location and function of the human frontal eye-field: a selective review. *Neuropsychologia* **34**, 475–483 (1996).
- Kastner, S. *et al.* Topographic maps in human frontal cortex revealed in memory-guided saccade and spatial working-memory tasks. *J. Neurophysiol.* **97**, 3494–3507 (2007).
- Fischer, J. & Whitney, D. Attention narrows position tuning of population responses in V1. *Curr. Biol.* **19**, 1356–1361 (2009).

ONLINE METHODS

Participants. Ten neurologically healthy volunteers (five female, 25 ± 2.11 years of age (mean \pm s.d.)) with normal or corrected-to-normal vision were recruited from the University of California, San Diego (UCSD). All participants provided written informed consent in accordance with the human participants Institutional Review Board at UCSD and were monetarily compensated for their participation. For the original experiment, participants participated in two to three scanning sessions, each lasting 2 h. Data from two participants (one female) were excluded from the main analysis because of excessive head movement (AJ3) or unusually noisy reconstructions during attend fixation runs (AG3).

In the follow-up experiment in which behavioral performance was carefully controlled and IPS subregions were retinotopically mapped, four participants of our original cohort were scanned for an additional two sessions, each lasting 1.5–2 h.

Stimulus. Stimuli were rear projected on a screen (90-cm width) located 380 cm from the participant's eyes at the foot of the scanner table. The screen was viewed using a mirror attached to the headcoil.

We presented an identical stimulus sequence during all imaging runs while asking observers to perform several different tasks. Each trial began with the presentation of a small red dot (T1) that was presented for 500 ms, followed by a flickering circular checkerboard stimulus at full contrast (stimulus radius (r_{stim}) = 1.163° , 1.47 cycles per degree) that was presented for 3 s and then a probe stimulus (T2) that was identical to T1. A 2-s intertrial interval separated each trial (Fig. 2a). T1 was presented between 0.176° and 1.104° from the center of the checkerboard stimulus along a vector of random orientation (in polar coordinates, θ_1 was randomly chosen along the range 0° to 360° , and r_1 was uniformly sampled from the range 0.176° to 1.104°). This ensured that the location of T1 was not precisely predictive of the checkerboard location. In 50% of the trials, T2 was presented in the same location as T1, and in the remaining trials, T2 was presented between 0.176° and 1.104° from the center of the checkerboard along a vector oriented at least 90° from the vector along which T1 was plotted (r_2 was uniformly sampled from the range 0.176° to 1.104° , and θ_2 was randomly chosen by adding between 90° and 270° , uniformly sampled, to θ_1). Polar coordinates used the center of the checkerboard stimulus as the origin. During the working memory condition (see below), participants based their response on whether T1 and T2 were presented in the exact same spatial position.

The location of the checkerboard stimulus was chosen pseudorandomly in each trial from a grid of 36 potential stimulus locations spaced by 1.163° , or r_{stim} . The stimulus location grid was jittered by 0.827° diagonally either up and to the left or down and to the right in each run, allowing for an improved sampling of space. All figures were presented aligned to a common space by removing jitter (see below).

In each run, there were 6 trials (1 trial for each stimulus location) and 9 null trials in which participants passively fixated for the duration of a normal trial (6 s). We scanned participants for between four and six runs of each task, always ensuring each task was repeated an equal number of times.

Tasks. Participants performed one of three tasks during each functional run (Fig. 2c). During attend fixation runs, participants responded when they detected a brief contrast dimming of the fixation point (0.33 s), which occurred in 50% of trials. During attend stimulus runs, participants responded when they detected a brief contrast dimming of the flickering checkerboard stimulus (0.33 s), which occurred in 50% of trials. During spatial working memory runs, participants made a button press response to indicate whether T2 was in the same location or a different location as T1. Notably, all three events (T1, checkerboard and T2) occurred during all runs, ensuring that the sensory display remained identical and that we were measuring changes in spatial representations as a function of task demands rather than changes as a result of inconsistent visual stimulation. For the follow-up behavioral control experiment, we dynamically adjusted the difficulty (contrast dimming or T1-T2 separation distance) to achieve a consistent accuracy of $\sim 75\%$ across tasks.

Eye tracking. Participants were instructed to maintain fixation during all runs. Fixation was monitored during scanning for four participants using an ASL LRO-R long-range eye-tracking system (Applied Science Laboratories) with a sampling rate of 240 Hz. We recorded mean gaze as a function of stimulus location and task

demands after excluding any samples in which neither the pupil nor the corneal reflection were reliably detected (Supplementary Fig. 1).

Imaging. We scanned all participants on a 3T GE MR750 research-dedicated scanner at UCSD. Functional images were collected using a gradient echo planar imaging (EPI) pulse sequence and an eight-channel head coil (19.2×19.2 cm FOV, 96×96 matrix size, 31 3-mm-thick slices with 0-mm gap, TR = 2,250 ms, TE = 30 ms, flip angle = 90°), which yielded a voxel size of $2 \times 2 \times 3$ mm. We acquired oblique slices with coverage extending from the superior portion of parietal cortex to the ventral occipital cortex.

We also acquired a high-resolution anatomical scan (FSPGR T1-weighted sequence, TR/TE = 11/3.3 ms, TI = 1,100 ms, 172 slices, flip angle = 18° , 1 mm³ resolution). Functional images were co-registered to this scan. Images were pre-processed using FSL (Oxford, UK) and BrainVoyager 2.3 (BrainInnovations). Preprocessing included unwarping the EPI images using routines provided by FSL, slice-time correction, three-dimensional motion correction (six-parameter affine transform), temporal high-pass filtering (to remove first-, second- and third-order drift), transformation to Talairach space and normalization of signal amplitudes by converting to z scores. We did not perform any spatial smoothing beyond the smoothing introduced by resampling during the co-registration of the functional images, motion correction and transformation to Talairach space. When mapping IPS subregions, we scanned the participants using an identical pulse sequence but instead used a 32-channel Nova Medical headcoil.

Functional localizers. All the ROIs used were identified using independent localizer runs acquired across multiple scanning sessions.

Early visual areas were defined using standard retinotopic procedures^{51,52}. We identified the horizontal and vertical meridians using functional data projected onto gray- and white-matter boundary surface reconstructions for each hemisphere. Using these meridians, we defined the areas V1, V2v, V3v, hV4, V2d and V3d. Unless otherwise indicated, data were concatenated across hemispheres and across the dorsal and ventral aspects of each respective visual area. We scanned each participant for between two and four retinotopic mapping runs ($n = 3$ completed two runs, $n = 3$ completed three runs, and $n = 2$ completed four runs).

hMT+ was defined using a functional localizer in which a field of dots either moved with 100% coherence in a pseudorandomly selected direction or were randomly replotted on each frame to produce a visual 'snow' display^{53,54}. Dots were each 0.081° in diameter and were presented in an annulus of between 0.63° and 2.26° around the fixation. During coherent dot motion, all dots moved at a constant velocity of $2.71^\circ \text{ s}^{-1}$. Participants attended the dot display for transient changes in velocity (during coherent motion) or replotted frequency (snow). Participants completed between one and three runs of this localizer ($n = 2$ completed one run, $n = 3$ completed two runs, and $n = 3$ completed three runs).

IPS and sPCS ROIs were defined using a functional localizer that required maintenance of a spatial location in working memory, a task that is commonly used to isolate IPS and sPCS, which is the putative human FEF^{47,49}. A flickering checkerboard subtending half of the visual field appeared for 12 s, during which time two spatial working memory trials were presented. During the flickering checkerboard presentation, we presented a red target dot for 500 ms, which was followed 2 s later by a green probe dot for 500 ms. After the probe dot appeared, participants indicated whether the probe dot was in the same location or a different location as the red target dot. Here we limited our definition of IPS to the posterior aspect (Supplementary Table 1). ROIs were functionally defined with a threshold of FDR-corrected $P < 0.05$ or a more stringent threshold when patches of activation abutted one another. Participants completed between one ($n = 2$) and two ($n = 6$) runs of this scan.

We also used data from these IPS and sPCS localizer scans to identify voxels in all other ROIs that were responsive to the portion of the visual field in which stimuli were presented in the main tasks, as the large checkerboard stimuli subtended the same visual area as the stimulus array used in the main task. All ROIs were masked on a participant-by-participant basis such that further analyses only included voxels with significant responses during this localizer task (FDR corrected $P < 0.05$).

Mapping IPS subregions. To determine the likely relative contributions of different IPS subregions to the localized ROI measured for all participants, we scanned the four participants who made up the behaviorally controlled cohort presented

in **Figures 6** and **7** using a polar angle mapping stimulus and an attentionally demanding task.

We used two stimulus types and behavioral tasks to define the borders between IPS subregions^{31,32,43,44}. In all runs, we used a wedge stimulus spanning a 72° polar angle and presented between 1.75° and 8.75° eccentricity rotating with a period of 24.75 s. In alternating runs, the wedge was either a 4-Hz flickering checkerboard stimulus (black-white, red-green or blue-yellow) or a field of moving black dots (0.3°, 13 dots per square degree², moving at 5° s⁻¹, changing direction every 8 s). During checkerboard runs, participants quickly responded after detecting a brief (250 ms) contrast dimming of a portion of the checkerboard. During moving dots runs, participants quickly responded after detecting a brief (417 ms) increase in dot speed. Targets appeared with 20% probability every 1.5 s. Difficulty was adjusted to achieve approximately 75% correct performance by changing the magnitude of the contrast dimming (checkerboard) or dot speed increment (moving dots) between runs. On average, participants performed with 84.1% accuracy in the contrast dimming task and 75.4% accuracy in the moving dots task. Two participants completed 14 runs (8 clockwise and 6 counterclockwise), and one participant completed 10 runs (AC; 5 clockwise and 5 counterclockwise). One participant was scanned with two different stimulus setups: half of all runs used the parameters described above and half used a wedge that spanned 60° of polar angle and rotated with a period of 36.00 s (AB; 6 runs clockwise, 6 runs counterclockwise).

Preprocessing procedures were identical to those used for the main task. To compute the best visual field angle for each voxel in IPS, we shifted the signals from counterclockwise runs earlier in time by twice the estimated hemodynamic response function (HRF) delay ($2 \times 6.75 \text{ s} = 13.5 \text{ s}$), removed the first and last full cycle of data (we removed 22 TRs for all participants except AB, for which we removed 32 TRs) and then reversed the time series so that all runs correspond to the clockwise stimulus presentation. We then averaged these time-inverted counterclockwise runs with the clockwise runs. We computed the power and phase at the stimulus frequency (1/24.75 Hz or 1/36 Hz, participant AB) and subtracted the estimated HRF delay (6.75 s) to align the signal phase in each voxel with the visual stimulus position. We then projected maps onto the reconstructed cortical surfaces for each subject and defined IPS0–IPS3 by identifying the upper and lower vertical meridian responses (**Supplementary Fig. 7a**). Low statistical thresholds were used (computed using normalized power at the stimulus frequency) to identify the borders of IPS subregions. Voxels were selected for further analysis by thresholding their activation during the same independent localizer task that was used to functionally define IPS and sPCS.

Encoding model. To measure changes in spatial representations under different task demands, we implemented an encoding model to reconstruct spatial representations of each stimulus used in the main task^{34,36–38}. This technique assumes that the signal measured in each voxel can be modeled as the weighted sum of different discrete neural populations, or information channels, that have different tuning properties³⁶. Using an independent set of training data, we estimated weights that approximate the degree to which each underlying neural population contributed to the observed BOLD response in each voxel (**Fig. 3a**). Next, an independent set of test data was used to estimate the activation within these information channels on the basis of the activation pattern across all voxels within an ROI in each test trial using the information channel weights in each voxel that were estimated during the training phase (**Fig. 3b**).

This approach requires specifying an explicit model for how neural populations encode information. Here we assumed a simple model for visual encoding within each ROI that focused exclusively on the spatial selectivity of visually responsive neural populations. To this end, we built a basis set of 36 two-dimensional spatial filters. We modeled these filters as cosine functions raised to a high power: $f(r) = (0.5 \cos(r\pi/s) + 0.5)^7$ for $r < s$ and 0 elsewhere (r is the distance from that filter's center; **Supplementary Fig. 2**). This allowed the filters to maintain an approximately Gaussian shape while reaching 0 at a fixed distance from the center (s°), which helped constrain curve-fitting solutions (below). The s (size constant) parameter was fixed at $5r_{\text{stim}}$, which is 5.8153°. The 36 identical filters formed a six-by-six grid spanning the visual space subtended by the stimuli. Filters were separated by 2.094°, with the centers tiled uniformly from 5.234° above, below, left and right of the fixation (**Fig. 3a**). The full-width half-maximum (FWHM) of all filters was 2.3103° (**Supplementary Figs. 2 and 3**). This ratio of filter size to spacing was chosen to avoid high correlations between

predicted channel responses (caused by too much overlap between channels, which can result in a rank-deficient design matrix) and to accomplish smooth reconstructions (if filters are too small, reconstructed spatial representations are 'patchy'; **Supplementary Fig. 3** shows an illustration of reconstruction smoothness as a function of the filter size to spacing ratio). All filters were assigned identical FWHMs so that known properties of the visual system, such as increasing receptive field size with eccentricity and along the visual stream^{39–41}, could be recovered without being built in to the analysis.

To avoid circularity in our analysis, we used a cross-validation approach to compute channel responses for every trial. First we used all runs but three (one run for each task condition) to create a training set that had an equal number of trials in each condition. Using this training set, we estimated channel weights within each voxel across all task conditions (i.e., runs one through five of the attend fixation, attend stimulus and spatial working memory tasks were used together to estimate channel weights, which were used to compute channel responses for run six of each task condition). The use of an equal number of trials from each condition in the training set ensures that channel weight estimation is not biased by any changes in BOLD response across task demands. Next the weights estimated across all task demand conditions were used to compute channel response amplitudes for each trial individually. Trials were then sorted according to their task condition and spatial location.

During the training phase, we created a design matrix that contained the predicted channel response for all 36 channels in every trial (**Fig. 3a**). The predicted response for each channel was computed by filtering a mask over the portion of the display subtended by the stimulus on that trial by the channel's basis function. The resulting design matrix was normalized to 1, such that reconstruction amplitudes are in units of BOLD z scores.

To extract relevant portions of the BOLD signal in every trial for computing channel responses, we took an average of the signal over two TRs beginning 6.75 s after trial onset. This range was chosen by examination of BOLD HRFs and was the same across all participants. Qualitatively, results do not change when other reasonable HRF lags are used, such as using two TRs starting 4.5 s after the stimulus.

Using this approach, we modeled voxel BOLD responses as a weighted sum of channel responses comprising each voxel^{36,38}. This can be written as a general linear model of the form

$$B_1 = WC_1 \quad (1)$$

where B_1 is the BOLD response in each voxel measured during every trial (m voxels $\times n$ trials), W is a matrix that maps channel space to voxel space (m voxels $\times k$ channels), and C_1 is a design matrix of predicted channel responses in each trial (k channels $\times n$ trials). The weight matrix \hat{W} was estimated by

$$\hat{W} = B_1 C_1^T (C_1 C_1^T)^{-1} \quad (2)$$

Then, using data from the held out test data set (B_2), the weight matrix estimated above was used to compute channel responses for every trial (\hat{C}_2), which were then sorted by task condition and spatial position:

$$\hat{C}_2 = (\hat{W}^T \hat{W})^{-1} \hat{W}^T B_2 \quad (3)$$

Reconstructing spatial representations. To reconstruct the region-wide representation of the visual stimulus viewed in every trial, we computed a weighted sum of the basis set using each channel response as the weight for the corresponding basis function (**Fig. 3b**). Reconstructions were computed out to a 5.234° eccentricity across the horizontal and vertical meridians, although visual stimuli only subtended at a maximum 4.523° eccentricity across the horizontal or vertical meridians. This was done to avoid edge artifacts in the reconstructions. Additionally, at this stage, the reconstructed visual fields were shifted to account for the slight jitter introduced in the presented stimulus locations and to align reconstructions from all trials. Runs in which stimuli were jittered up and to the left were reconstructed by moving the centers of the basis functions down and to the right, and runs in which stimuli were jittered down and to the right were reconstructed by moving the centers of the basis functions up and to the left. These shifts serve to counter the spatial jitter of stimulus presentation for visualization and quantification. By including spatial jitter during stimulus presentation, we are able to attain a more nuanced estimate of channel weights by sampling 72 stimulus locations rather than 36.

We averaged each participant's reconstructions at all 36 spatial locations for each task condition across trials. For **Figure 4**, all $n = 8$ participants' average reconstructions for each task condition were averaged, and reconstructions from all ROIs and task conditions were visualized on a common color scale to illustrate differences in spatial representations across the different task conditions and spatial locations. The three-by-three subset of reconstructions shown in **Figure 4b** was chosen as it is representative and results were similar for all quadrants.

For the follow-up control experiment, we plotted reconstructed spatial representations from only correct trials by co-registering all representations for trials at matching eccentricities and then averaging across all co-registered representations for each participant at each eccentricity. We co-registered representations for like eccentricities to the top left quadrant (**Fig. 6b**). Representations were rotated in 90° steps and flipped across the diagonal (equivalent to a matrix transpose operation on pixel values) as necessary.

Notably, this analysis depends on two necessary conditions. First, individual voxels must respond to certain spatial positions more than others, although the shape of these spatial selectivity profiles is not constrained to follow any particular distribution (for example, it need not resemble a Gaussian distribution). Second, the spatial selectivity profile for each voxel must be stable across time, such that spatial selectivity estimated on the basis of data in the training set can generalize to the held-out test set.

Curve fitting. To quantify the effects of attention on visual field reconstructions, we fit a basis function to all 36 average reconstructions for each participant for each task condition for each ROI using `fminsearch` as implemented in MATLAB 2012b (which uses the Nelder-Mead simplex search method; Mathworks, Inc).

The error function used for fitting was the sum of squared errors between the reconstructed visual stimulus and the function

$$f(r) = b + a \times (0.5 + 0.5 \cos(r\pi/s))^7 \quad \text{for } r < s, 0 \text{ elsewhere} \quad (4)$$

where r is computed as the Euclidean distance from the center of the fit function. We allowed the baseline (b), amplitude (a), location (x, y) and size (s) to vary as free parameters. The size s was restricted so as not to be too large or too small (confined to $0.5815^\circ < s < 26.17^\circ$), and the location was restricted around the region of visual stimulation (x and y lie within stimulus extent borders $+1.36^\circ$ on each side).

Because of the number of free parameters in this function, we performed a two-step stochastic curve-fitting procedure to find the approximate best-fit function for each reconstructed stimulus. First we averaged reconstructions for each spatial location across all three task conditions and performed 50 fits with random starting points. The fit with the smallest sum squared error was used as the starting point around which all other starting points were randomly drawn when fitting to reconstructions from each task condition individually. When fitting individual task condition reconstructions, we performed 150 fits for each condition. We used parameters from the fit with the smallest sum squared error as a quantitative characterization of the reconstructed visual stimulus. Then we averaged the fit parameters across like eccentricities within each task condition, ROI and participant. For the follow-up control experiment, we performed an identical fitting procedure on each of the co-registered representations to directly estimate the best fit parameters at each eccentricity.

Excluded participant. For one participant (AG3), reconstructions from the attend fixation runs were unusually noisy and could not be well approximated by the basis function used for fitting. However, both the attend stimulus and spatial working memory runs for this individual exhibited successful reconstructions (**Supplementary Fig. 4**). As the estimated channel weights used to compute these stimulus reconstructions were identical across the three task conditions, only changes in information coding across task demands could account for this radical shift in reconstruction fidelity. Because this participant's reconstructions could not be accurately quantified for the attend fixation condition, the reconstructions and fit parameters for this individual for all conditions have been left out of the data presented in the Results. However, as noted above, data from this participant are consistent with our main conclusion that attentional demands influence the quality of spatial representations.

Evaluating the relationship between amplitude and size. It may be the case that our observation of increasing spatial representation size with increasing stimulus eccentricity is purely a result of intertrial variability in the reconstructed stimulus position. That is, the same representation could be jittered across trials, and the resulting average representation across trials would appear 'smeared' and would be fit with a larger size and smaller amplitude. If this were true, changes in these parameters would always be negatively correlated with one another—an increase in size across conditions would always occur with a decrease in amplitude.

To evaluate this possibility, for each eccentricity, ROI and condition pair (attend stimulus and attend fixation, spatial working memory and attend stimulus, and spatial working memory and attend fixation), we correlated the change in size with the change in amplitude (each correlation contained eight observations, corresponding to $n = 8$ participants). To evaluate the statistical significance of these correlations, we repeated this procedure 10,000 times, each time shuffling the condition labels separately for size and amplitude, recomputing the difference and then recomputing the correlation between changes in size and changes in amplitude. This resulted in a null distribution of chance correlation values against which we determined the probability of obtaining the true correlation value by chance. After correction for FDR, no correlations were significant (all $P > 0.05$; of note, FDR is more liberal than Bonferroni correction).

Representations from the ventral and dorsal aspects of V2 and V3. For **Supplementary Figure 5a**, we generated reconstructions using a procedure identical to that used for **Figure 4**, except we only used voxels that were assigned to the dorsal or ventral aspects of V2 and V3 instead of combining voxels across the dorsal and ventral aspects, as was done in the main analysis.

Reconstructions of untrained (new) stimuli. For **Supplementary Figure 5b**, we estimated channel weights using all runs of all task conditions from the main task as a training set. We used these weights to estimate channel responses from the BOLD data taken from an entirely new data set, which consisted of responses to a hemi-annulus-shaped radial checkerboard (**Supplementary Fig. 5b**).

This new experiment featured four stimulus conditions: left-in, left-out, right-in and right-out. The inner hemi-annuli subtended at 0.633° to 2.262° eccentricity. The outer hemi-annuli subtended at 2.262° to 4.523° eccentricity. Stimuli were flickered at 6 Hz for 12 s in each trial while the participants performed a spatial working memory task on small probe stimuli presented at different points within the displayed stimulus.

The BOLD signal used for reconstruction was taken as the average of four TRs beginning 4.5 s after stimulus onset. These data were used as the test set. Otherwise, the reconstruction process was identical to that in the main experiment, as were all other scan parameters and preprocessing steps.

pRF estimation. To determine whether the spatial sensitivity of each voxel across all trials and all runs changed across conditions, we implemented a new version of a pRF analysis^{42,55}. For this analysis, we estimated the unimodal, isotropic pRF that best accounts for the BOLD responses to each stimulus position within every single voxel. This analysis is complementary to the primary analyses described above.

For four participants (those presented in **Figs. 6** and **7** and **Supplementary Fig. 7**) and four ROIs for each participant (V1, hV4, hMT+ and IPS0, which were chosen because this set includes both ROIs with (hV4, hMT+ and IPS0) and without (V1) attentional modulation), we used data across all runs within each task condition and ridge regression⁵⁶ to identify pRFs for each voxel under each task condition. We computed these pRFs using a method similar to that used to compute channel weights in the encoding model analysis (**Fig. 3a** and Online Methods; the univariate step 1 of the encoding model, see equation (1)). We generated predicted responses with the same information channels that were used for the encoding model analysis (**Fig. 3a**), and reconstructed pRFs for each task condition for a given voxel were defined as the corresponding spatial filters weighted by the computed weight for each channel (**Supplementary Fig. 8a**).

In the main analysis in which we computed spatial reconstructions on the basis of activation patterns across an entire ROI (**Figs. 4** and **6c**), any spatial information encoded by a voxel's response could be exploited; this is true even if the voxel's response to different locations was not unimodal (it need not follow any set distribution, as long as it responds consistently). However, univariate pRFs

computed on a voxel-by-voxel basis cannot be well characterized by an isotropic function if they are not unimodal⁵⁷. Thus, to ensure that most of the pRFs were sufficiently unimodal to fit an isotropic function, we used ridge regression^{56,57} when computing spatial filter weights for the pRF analysis. The regression equation for computing channel weights then becomes

$$\hat{W}^T = (C_1 C_1^T + \lambda I)^{-1} C_1 B_1^T \quad (5)$$

where I is an identity matrix ($k \times k$). To identify an optimal ridge parameter (λ), we computed the Bayes information criterion⁵⁸ value across a range of λ values (0 to 500) for each voxel using data concatenated across all three task conditions. This allowed for an unbiased selection of λ with respect to task condition. The λ with the minimum mean Bayes information criterion value across all voxels within a ROI was selected, and this λ was used to compute channel weights for each of the three task conditions separately. An increasing λ value results in greater sparseness of the best-fit channel weights for each voxel, and a λ value of 0 corresponds to ordinary least-squares regression.

After computing pRFs for each task condition, we fit each pRF with the same function that was used to fit the spatial representations (equation (4)) using a similar optimization procedure. We restricted the fit size (FWHM) to be at most 8.08°, which corresponds to nearly the full diagonal distance across the stimulated visual field. This boundary was typically encountered only for hMT+ and IPS0 and served to discourage the optimization procedure from fitting large, flat surfaces. Then we computed an R^2 value for each fit and used only voxels for which the minimum R^2 across conditions was greater than or equal to the median of the minimum R^2 across conditions from all voxels in that participant's ROI (Supplementary Fig. 8a,b).

Because we only have a single parameter estimate for each condition for each voxel, we evaluated whether fit size is more likely to increase or decrease between each pair of task conditions (attend stimulus compared to attend fixation, spatial working memory compared to attend stimulus and spatial working memory compared to attend fixation) for each region for each participant by determining the percentage of voxels that lie above the unity line in a plot of one condition against another (Supplementary Fig. 8d).

Simulating data with different pRF properties. In order to assess whether our region-level multivariate spatial representation analysis would be sensitive to changes in voxel-level univariate pRFs, we generated simulated data using two different pRF modulation models.

For the first model (Supplementary Fig. 10a,b), we randomly generated 500 pRF functions so as to uniformly sample the visual field for each of two conditions (condition A, smaller pRFs; condition B, larger pRFs). Across the two conditions, each simulated voxel's pRF maintained its preferred position while its amplitude and baseline were each randomly and independently sampled across conditions from the same normal distribution (amplitude: $\mu = 0.8513$, $\sigma = 0.25$; baseline: $\mu = -0.1952$, $\sigma = 0.25$; these values were taken from the average-fit pRF parameters across all participants for hV4 in the attend fixation and attend stimulus conditions; Supplementary Fig. 9a). pRF size (FWHM) was sampled from a normal distribution with $\sigma = 0.5$ and a mean of $\mu = 4.405^\circ$ for condition A (mean of pRF size for hV4, attend fixation) and $\mu = 4.89^\circ$ for condition B (mean of pRF size for hV4, attend stimulus; an increase of 11%). In our simulation, this resulted in 79% of simulated voxels showing larger pRF sizes in condition B compared to condition A. For the second model (Supplementary Fig. 10c,d), we used the upper median split of fit pRFs for the single participant shown in Supplementary Figure 8c, hV4 ROI, to generate the simulated BOLD data. This allowed us to simulate region-level BOLD data for each attention condition tested in our experiment and enabled us to determine whether the changes in univariate voxel-level pRF size we observed (Supplementary Fig. 9) are consistent with the multivariate region-level spatial representations presented in the main text (Figs. 5 and 7).

After generating voxel-level pRFs using each of the two models described above, we added noise to the simulated weights (Gaussian noise added independently to each channel weight, $\sigma = 0.1$) and presented model voxels with six runs of all 36 spatial positions for each condition. We simulated each voxel's BOLD response as the predicted channel response (response of corresponding spatial filter; Fig. 3a) to each stimulus weighted by the corresponding channel weights. We added Gaussian noise to the resulting BOLD data for each simulated voxel independently ($\sigma = 0.1$). Then all analyses of multivariate spatial representations

proceeded identically to those described above. We computed spatial representations using estimated channel weights computed across all conditions within a model (i.e., condition A and condition B (Supplementary Fig. 10a,b) or attend fixation, attend stimulus and spatial working memory (Supplementary Fig. 10c,d)) and then fit the average spatial representations with a smooth surface (Online Methods) to determine the amplitude and size of each spatial representation. We then averaged these parameters across all 36 positions.

Statistical procedures. All behavioral analyses on accuracy data were performed using two-way repeated-measures ANOVA, with task condition and stimulus eccentricity modeled as fixed effects (three levels and six levels, respectively; Figs. 2d and 6a).

To assess whether fit parameters to reconstructed spatial representations reliably changed as a function of task demands, we performed a multistage permutation testing procedure. This nonparametric procedure was adopted because the spatial filters (basis functions) used to estimate the spatial selectivity of each voxel during the training phase (Fig. 3a) overlapped and were not independent (violating a key assumption of standard statistical tests).

For each parameter (the rows in Figs. 5 and 7), we first found ROI-parameter combinations that showed an omnibus main effect in a repeated-measures ANOVA (1 factor, 18 levels) corrected using a FDR algorithm⁵⁹ across all ROIs. Then we computed F scores for a two-way repeated measures design with eccentricity and condition as factors (six levels and three levels, respectively) for ROIs with significant omnibus main effects.

For all tests, because we had a relatively small n ($n = 8$ for Fig. 5, and $n = 4$ for Fig. 7 and Supplementary Fig. 7) and the range of parameters was in some cases restricted to be positive (size), we computed an F distribution for the null hypothesis that there is no main effect of the omnibus factor (omnibus test) or that there is no main effect of condition, eccentricity or their interaction (for a follow-up two-way test) by shuffling trial labels within each participant 100,000 times. For each data permutation, we computed a new F score for the omnibus test, and for ROI-parameter combinations with a significant omnibus effect, we computed a main effect of condition, eccentricity and their interaction. P values were estimated as the probability that the F score computed based on the shuffled data was equal to or greater than the F scores computed using the actual data. These additional tests were corrected for multiple comparisons using Bonferroni's method within each parameter. We also occasionally highlight trends in the data by reporting P values that did not reach significance under correction for multiple comparisons at this sample size as marginal effects, and such P values are always reported as being uncorrected in the text. For display purposes, marginally significant tests are shown in Figures 5 and 7 at uncorrected $P < 0.025$.

In addition, we performed a three-factor repeated-measures ANOVA with ROI, task condition and eccentricity modeled as fixed effects to determine whether the fit parameters changed across ROIs ($n = 8$ for Fig. 5, and $n = 4$ for Fig. 7). We implemented the same permutation procedure described above to compute P values (10,000 iterations).

To determine whether pRF size increases at higher eccentricities, we computed a linear fit to a plot of each voxel's pRF size compared to its pRF eccentricity for each ROI for each condition for each participant (Supplementary Fig. 8c). To determine whether the slope of the fit line was reliably positive for a given ROI, participant and condition, we computed confidence intervals around the best-fit slopes using bootstrapping (resampled all voxels with replacement 10,000 times), and the related P value was defined as the as the probability that the slope was ≤ 0 . We used a Bonferroni-corrected significance threshold for 48 planned comparisons (4 ROIs \times 4 participants \times 3 conditions) of $\alpha = 0.001$ (Supplementary Results and Supplementary Table 2).

To evaluate the statistical significance of the pRF size increase (Supplementary Fig. 9), we first performed a two-way repeated-measures ANOVA with ROI and condition modeled as fixed effects and participant modeled as a random effect in which we shuffled ROI and condition labels for each participant and recomputed the percentage of voxels that increased in size across each condition pair. We repeated this shuffling procedure 10,000 times and compared F scores computed using the real labels to the distribution generated using the shuffled labels, as described above. Then we compared whether each condition pairing resulted in a significant change in pRF size for each ROI by computing a T score testing against the null hypothesis that 50% of voxels show an increase in pRF size. As described above, we generated a null T distribution by shuffling condition labels

within each participant 10,000 times. For this analysis, we used a Bonferroni-corrected significance threshold for 12 planned comparisons (4 ROIs \times 3 conditions) of $\alpha = 0.0042$.

51. Engel, S.A. *et al.* fMRI of human visual cortex. *Nature* **369**, 525 (1994).
52. Sereno, M.I. *et al.* Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Science* **268**, 889–893 (1995).
53. Tootell, R.B. *et al.* Functional analysis of human MT and related visual cortical areas using magnetic resonance imaging. *J. Neurosci.* **15**, 3215–3230 (1995).
54. Serences, J.T. & Boynton, G.M. The representation of behavioral choice for motion in human visual cortex. *J. Neurosci.* **27**, 12893–12899 (2007).
55. Kay, K.N., Naselaris, T., Prenger, R.J. & Gallant, J.L. Identifying natural images from human brain activity. *Nature* **452**, 352–355 (2008).
56. Hoerl, A.E. & Kennard, R.W. Ridge regression: biased estimation for nonorthogonal problems. *Technometrics* **12**, 55–67 (1970).
57. Lee, S., Papanikolaou, A., Logothetis, N.K., Smirnakis, S.M. & Keliris, G.A. A new method for estimating population receptive field topography in visual cortex. *Neuroimage* **81**, 144–157 (2013).
58. Schwarz, G. Estimating the dimension of a model. *Ann. Stat.* **6**, 461–464 (1978).
59. Benjamini, Y. & Yekutieli, D. The control of the false discovery rate in multiple testing under dependency. *Ann. Stat.* **29**, 1165–1188 (2001).