ATTENTION MODULATES STIMULUS REPRESENTATIONS IN NEURAL FEATURE DIMENSION MAPS

Daniel D. Thayer (danielthayer@ucsb.edu) and **Thomas C. Sprague** (tsprague@ucsb.edu)

Department of Psychological and Brain Sciences, University of California, Santa Barbara

Running head: Attention sculpts feature dimension maps

Key words: feature-selective attention, visual system, fMRI, priority map

Word count: 8558

Acknowledgments: This work was supported by NIH R01-EY035300, an Alfred P Sloan Research Fellowship, an Nvidia Hardware Grant, and US Army Research Office Cooperative Agreement W911NF-19-2-0026 for the Institute for Collaborative Biotechnologies

ABSTRACT

Computational theories posit that attention is guided by a combination of spatial maps for individual features that can be dynamically weighted according to task goals. Consistent with this framework, when a stimulus contains several features, attending to one or another feature results in stronger fMRI responses in regions preferring the attended feature. We hypothesized that multivariate activation patterns across feature-responsive cortical regions form spatial 'feature dimension maps', which combine to guide attentional priority. We tested this prediction by reconstructing neural feature dimension maps from fMRI activation patterns across retinotopic regions of visual cortex using a spatial inverted encoding model. Participants viewed a peripheral visual stimulus at a random location which always contained moving colored dots. On each trial, participants were precued to report the predominant direction of motion or color of the stimulus, or to attend fixation. Stimulus representations in reconstructed priority maps were selectively enhanced in color-selective regions when color was attended, and in motion-selective regions when motion was attended, and this effect was primarily observed at the stimulus location. While enhancement was localized to the stimulus in color-selective regions, motion-selective regions were globally enhanced when motion was task relevant. These results suggest different cortical regions support spatial maps of different visual features, and that each map is uniquely reweighted based on task demands to guide visual behavior towards the most relevant locations based on important features.

Introduction

The environment typically contains many objects competing for our attention. While gardening, your goal may be to find green weeds in a strawberry bush, so attention will be directed to green objects. However, task-irrelevant objects that are salient may capture your attention, such as a hummingbird darting in the sky or a red strawberry. The factors which contribute to how attention is allocated in the visual field—task-relevant goals and task-irrelevant salience—are described by priority map theory (Awh et al., 2012; Itti & Koch, 2001; Luck et al., 2020; Serences & Yantis, 2006; Treisman & Gelade, 1980; Wolfe, 1994), where these priority signals are integrated to index the most important locations in a scene. Locations with the highest level of priority are then selected for further processing (Carrasco, 2011; Eckstein, 2011).

Per this theory, overall priority is computed based on the combination of priority indexed independently for individual feature dimensions, like color, motion, orientation, and shape. Priority within these 'feature dimension maps' is a combination of salience (defined based on image-computable properties, like local feature contrast) and relevance (defined based on task goals, like which location(s) or feature(s) are important; Wolfe & Horowitz, 2004, 2017). For example, a color dimension map would index locations made salient based on their unique color relative to other stimuli, while a motion dimension map would do the same for stimuli moving a unique direction. Additionally, it would be expected that making one feature task-relevant would sculpt stimulus representations in the corresponding map, leading to a greater impact on downstream priority maps (Bahle et al., 2019; Corbetta & Shulman, 2002; Folk et al., 1992; Leber & Egeth, 2006; Moore & Armstrong, 2003; Poltoratski et al., 2017; Witkowski & Geng, 2023).

Recently, we demonstrated that color- (hV4/VO1/VO2; Bartels & Zeki, 2000; Brewer et al., 2005) and motion- (TO1/TO2; Amano et al., 2009; Huk et al., 2002) selective visual cortical regions index salient locations on the screen even when those locations are task-irrelevant based on their preferred feature values (Thayer & Sprague, 2023). However, how activation profiles across these maps index *task relevance* of cued feature dimensions remains unknown.

Previous work has shown that feature-selective attention modulates fMRI BOLD activity in color and motion regions (Beauchamp et al., 1997; Chawla et al., 1999; McMains et al., 2007; O'Craven et al., 1999; Runeson et al., 2013; Saenz et al., 2002), which provides some support that feature-selective areas are modulated by attention to specific features. Generally, these results show that attending to a specific feature value (e.g., green) enhances BOLD signals corresponding to stimuli matching the attended feature value across the entire screen (Saenz et al., 2002), and that attending to one feature dimension of a multi-feature object leads to enhanced BOLD responses in the corresponding feature-selective region (McMains et al., 2007; Runeson et al., 2013). While aspects of the impact of attention to stimulus features on neural representations have been characterized, how attention to a specific feature dimension sculpts precise spatial profiles of activation across these regions is less clear. Specifically, it is unknown whether feature-selective modulations occur equivalently across the entire visual field within feature-selective visual regions, reflecting a *global enhancement* of processing of the attended feature dimension across the entire map (Chawla et al., 1999; Saenz et al., 2002; Serences & Boynton, 2007; Treue & Trujillo, 1999), or instead, are restricted to the location of the presented stimuli containing the attended feature dimension, reflecting a *local enhancement* of relevant locations within the attended map (Fig. 1) (McMains et al., 2007).

Here, we used a multivariate inverted encoding model (IEM; Thayer & Sprague, 2023) approach to determine how attention to a cued feature dimension modulated activity across neural feature dimension maps encoding an identical visual stimulus. We found that map locations corresponding to the stimulus location in both regions had the greatest activation, and that these stimulus representations were strongest when attending to each region's preferred feature. Interestingly, feature-specific modulations in motion regions showed evidence for global enhancement while enhancement in color regions was localized to the stimulus, meaning that each feature dimension was uniquely modulated by feature-selective attention. These findings were only observable after implementing IEM and leveraging region-level spatial activation patterns, demonstrating the necessity to consider the information jointly encoded across entire neural populations when characterizing the computational properties of neural feature dimension maps.

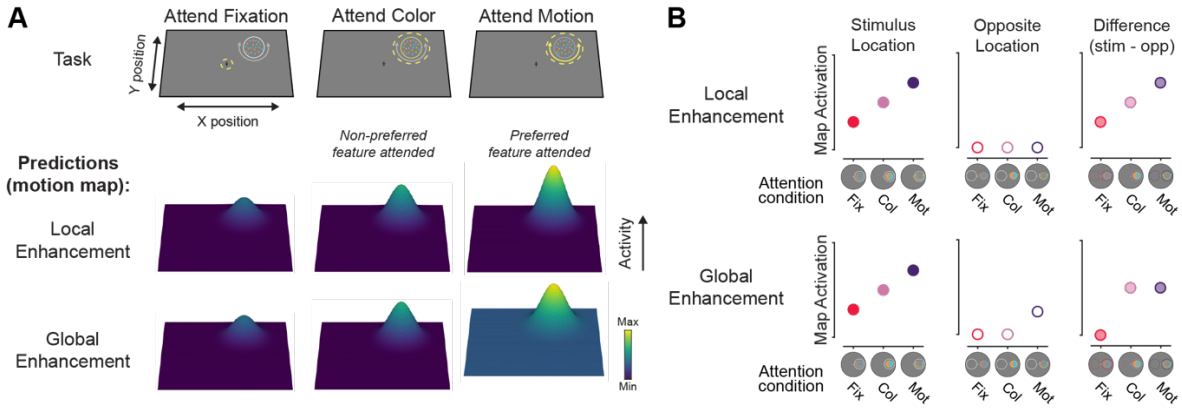

Figure 1: How does attention to a specific feature dimension sculpt neural feature dimension maps? **A.** Priority map theory posits that various "feature dimension maps" are used to index the most important locations in the visual field based on computations within their preferred feature dimension, and that activation in these maps should scale with an observer's goals. If detecting or discriminating motion is needed for an ongoing task (e.g., identifying the motion direction of a darting hummingbird), then activation within the corresponding "motion map" would increase the importance associated with the hummingbird's location. There are two ways in which the motion map could prioritize information beyond the local effects expected of spatial attention (e.g., Sprague et al, 2018). Local enhancement could occur, such that only the location of a stimulus with the attended feature is prioritized. Alternatively, global enhancement could occur, such that activation across the entire map is additively scaled, leading to increased sensitivity to the attended feature dimension at any location. This type of modulation would still drive a stronger target representation but would additionally lead to stronger responses in locations that contain no stimulus when motion is the goal-relevant feature dimension. A motion dimension map is depicted here, but modulations would similarly apply for other feature dimensions, such as color. **B.** Evaluating the activation at the stimulus location and the opposite location within the feature (motion) map can discriminate between the local and global enhancement accounts. Both models predict that the stimulus will have the greatest activation at the stimulus location when the preferred feature dimension is relevant (e.g., motion; left). If enhancement is localized, then activation at the opposite location should not change across conditions (center). However, if there is global enhancement, then the opposite location should have increased activation when motion is task relevant. By computing the difference in activation between the stimulus and opposite locations, the spatial specificity of feature-based modulations can be assessed (right). If this activation difference (stimulus-opposite) is greater for the attend motion condition in the motion map, enhancement is localized. However, if the activation difference is similar between the attend color and motion conditions, then enhancement is global across the feature dimension map.

Materials & Methods

Participants

We recruited 10 subjects from the University of California, Santa Barbara (UCSB) community to participate in this study (9 female, 18-29 years old). We identified an appropriate sample size for our main effect of interest by conducting a power analysis using pilot data (n = 3). In the pilot study, we had a strong effect (d_z = 1.24) which required a sample of 6 subjects to obtain 80%

power with an alpha criterion of α = .05 as indicated by a power analysis conducted in G*Power (Faul et al., 2007). We collected a large number of measurements from each subject to minimize within-subject variance, which can benefit statistical power more than increased sample sizes (Baker et al., 2021). Subjects reported normal or corrected-to-normal vision and did not report neurological conditions. All procedures were approved by the UCSB Institutional Review Board (IRB# 5-24-0030), and the study was registered on ClinicalTrials.gov (NCT06281457). Subjects gave written informed consent before participating and were compensated for their time ($20/h for scanning sessions, $10-20/h for behavioral familiarization/training).

Stimuli and Procedure

Participants performed a 30-minute training session before scanning so that they were familiarized with the instructions. We used this session to establish the initial behavioral performance thresholds used in the first run of the scanning session. In the main task session, participants were scanned for a single two-hour session consisting of at least 4 runs of a spatial mapping task, which were used to independently estimate encoding models for each voxel, and 8 runs of the experimental feature-selective attention task. All participants also underwent additional anatomical and retinotopic mapping scanning sessions (1-2x 1.5-2 hr sessions) to identify regions of interest (ROI; see Region of interest definition).

Stimuli were presented using the Psychophysics toolbox (Brainard, 1997; Pelli, 1997) for MATLAB (The MathWorks, Natick, MA). Visual stimuli were rear-projected onto a screen placed ~110 cm from the participant's eyes at the head of the scanner bore using a contrast-linearized LCD projector (1,920×1,080, 60 Hz) during the scanning session. For two of the participants, a CRS BOLDScreen LCD monitor was used to present stimuli (1,920x1,080, 60 Hz) that was placed ~149 cm from the participant's eyes. Participants who performed the task on the BOLDScreen had slightly different stimulus dimensions (scaled by a factor of .9 compared to the projector), but all results are qualitatively identical for the subjects who used either the BOLDScreen monitor or the projector, and all findings hold considering only the n = 8 participants tested with the projector. Throughout, we describe stimulus dimensions/units based on the projector stimulus size. In the behavioral familiarization session, we presented stimuli on a contrast-linearized LCD monitor (2,560×1,440, 60 Hz) 62 cm from participants, who were seated in a dimmed room and positioned

using a chin rest. For all sessions and tasks (main tasks, localizers, and mapping task), we presented stimuli on a neutral gray circular aperture (9.15° radius), surrounded by black (only aperture shown in Fig. 2).

Feature-selective attention task

For each trial of the main task (Fig. 2), participants viewed a colorful moving dot stimulus in the periphery (similar to that used in Mante et al., 2013) along with a flickering central fixation cross (0.2° size). The colorful moving dot stimulus was a 1.5°-radius circular disc comprised of several individual dots centered 5° from fixation at a random location along an invisible ring. Individual dots were physically isoluminant, colored either cyan (RGB: [0; 200; 200]) or orange (RGB: [255; 165; 0]), and moved clockwise or counterclockwise around the stimulus center, or in a random direction. Individual dots occupied 0.04° of visual angle, with a dot density was 40 dots/deg^2, and moved at a speed of 4.5°/s. Dots were randomly replotted every 50 ms or when they exceeded the stimulus bounds. Both cyan and orange dots were always present. Dots were always moving, but only some dots on each trial would move clockwise or counterclockwise, while the rest of the dots moved in a random direction (uniform distribution of planar motion directions).

At the start of each trial, participants were shown a letter cue at fixation (0.4° height) that would indicate what aspect of the display to attend. The letter cue was presented for 1000 ms in white (RGB: [180; 180; 180]) Arial font. The cue could be an 'F' (attend fixation), 'C' (attend color), or 'M' (attend motion). After a 500 ms blank period, with only the aperture and fixation dot present, the colorful moving dot stimulus and fixation cross was shown for 3000 ms. When cued to attend the color of the stimulus ('C'), participants reported which individual feature value (cyan or orange) was most prevalent in the stimulus. When cued to attend the motion of the stimulus ('M'), they reported the coherent motion direction of the stimulus (clockwise or counterclockwise). Since motion and color was present on each trial, visual input was constant, and we were able to isolate changes in the neural stimulus representations across brain regions due to instructed task goals. We adjusted the difficulty of the attend color/motion tasks between runs by changing the coherence of the dots to try and equate task performance between conditions (e.g., 80% dot coherence indicates that 80% of the dots were presented in cyan or were moving clockwise). The

first run of the task in the scanner used coherence values acquired from the behavioral training session.

As a neutral baseline condition, participants attended the central fixation cross on some trials. The vertical and horizontal lines of the fixation cross were 0.2° of visual angle long (0.1° from center in either direction) and flickered at 3 Hz (10 frames on, 10 frames off at 60 Hz). On each trial, one line of the fixation cross would slightly change in size for a single flicker cycle. When cued to attend fixation ('F'), participants reported which line increased in length (horizontal or vertical) with a button press. Since the dot stimulus was ignored on these trials due to fixation task difficulty (Thayer & Sprague, 2023), the neural response associated with the stimulus on these trials was purely driven by visual input. We adjusted the difficulty of the fixation task between runs by altering the degree of size change for vertical/horizontal lines based on behavioral accuracy (range: 0.21° to 0.32°). The fixation target appeared at a random time during the stimulus presentation period other than the first or last 500 ms.

Each run contained 24 trials, with 8 trials of each condition (attend fixation, attend stimulus color, attend stimulus motion). All trials were separated by a randomly selected ITI ranging from 6-9 s with an average ITI of 7.5 s. Trial order was shuffled within run. Each run started with a 3 s blank period and ended with a 10.5 s blank period, for a total run duration of 301.5 s. Eye position (right eye) was monitored throughout the experiment using an Eyelink 1000 eyetracker (SR Research) at 500 Hz. Participants performed a 5-point calibration at the start of the experiment.

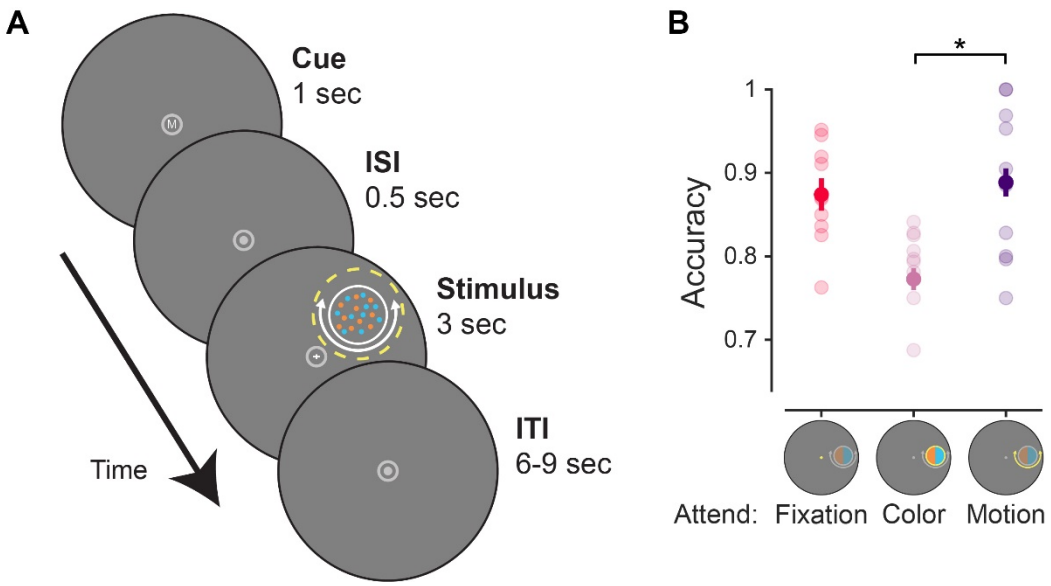

Figure 2: Feature-selective attention task. A. On each trial, participants viewed a dot array comprised of colorful moving dots. All dots were randomly either cyan or orange, and all dots were always moving, but only some would move clockwise or counterclockwise, while the rest moved in random direction. A flickering fixation cross (3 Hz) was also always present, and the vertical or horizontal arm of the cross would increase in size once during the stimulus presentation period. A letter at fixation cued participants to perform one of three attention tasks related to the upcoming stimulus display (F: attend fixation, C/M attend peripheral stimulus color/motion, respectively). On attend fixation trials, participants reported whether the horizontal or vertical arm of the fixation cross lengthened; on attend color/motion trials, they reported the prominent feature value of the cued feature dimension (more cyan/orange dots, or more dots moving clockwise/counterclockwise). **B.** Accuracy for each attention condition. Performance was slightly worse for the attend color than the attend motion or fixation conditions. Response times (not shown) were not statistically different between the attend color (mean RT: 1.66; SEM: 0.05) and attend motion (mean RT: 1.58; SEM 0.06) conditions (p = 0.27). Error bars indicate SEM across participants. * Significant difference based on permuted paired-samples t test (p < 0.05).

Spatial mapping task

We also acquired several runs of a spatial mapping task used to independently estimate a spatial encoding model for each voxel, following previous studies (Sprague et al., 2016; Sprague, Itthipuripat, et al., 2018; Thayer & Sprague, 2023). On each trial of the mapping task, we presented a flickering checkerboard at different positions selected from a hexagonal grid spanning the screen. Participants viewed these stimuli and responded whenever a rare contrast change occurred (6 out of 43 trials, 13.9%), evenly split between contrast increments and decrements. The checkerboard stimulus was the same size as the stimulus in the feature-selective attention task (1.5° radius) and was presented at 70% contrast and 6-Hz full-field flicker. All stimuli appeared within a gray circular aperture with a 9.15° radius, as in the feature-selective attention task. For each trial, the location of the stimulus was selected from a hexagonal grid of 37 possible locations

with an added random uniform circular jitter (0.5° radius). The base position of the grid was rotated by 30° on every other scanner run to increase spatial sampling density. As a result, every mapping trial was unique, which enabled robust spatial encoding model estimation.

Each trial started with a 3000 ms stimulus presentation period. If a target was present, then the stimulus would be dimmed/brightened for 500 ms with the stipulation that the contrast change would not occur in either the first or last 500 ms of the stimulus presentation period. Finally, ITIs were randomly selected to be either 6 or 8.25 s. All target-present trials were discarded when estimating the spatial encoding model. Each run consisted of 43 trials (6 of which included targets). We also included a 3 s blank period at the beginning of the run and a 10.5-s blank period at the end of the run. Each run totaled 432 s.

Retinotopic mapping task

We used a previously reported task (Mackey et al., 2017; Thayer & Sprague, 2023) to identify retinotopic regions of interest (ROIs) via the voxel receptive field (vRF) method (Dumoulin & Wandell, 2008). Each run of the retinotopy task required participants attend several random dot kinematograms (RDK) within bars that would sweep across the visual field in 2.25 s steps. Three equally sized bars were presented on each step and the participants reported in which of the two peripheral bars the motion in the central bar matched with a button press. Participants received feedback via a red (incorrect) or green (correct) color change at fixation. We used a three-down/one-up staircase to maintain ~80% accuracy throughout each run so that participants would continue to attend the RDK bars. RDK bars swept 17.5° (15.8° on BOLDScreen) of the visual field. Bar width and sweep direction was pseudo-randomly selected from several different widths (ranging from 2.0° to 7.5°) and four directions (left-to-right, right-to-left, bottom-to-top, and top-to-bottom).

fMRI acquisition and MRI preprocessing

All scanning was conducted at the UCSB Brain Imaging Center with a 3T Siemens Prisma scanner. fMRI data acquisition and preprocessing pipelines in the current study exactly follow our previous report (Thayer & Sprague, 2023). All fMRI scans for retinotopic mapping, model estimation, and main experimental task were acquired using the CMRR MultiBand accelerated EPI pulse

sequences. We acquired all images with the Siemens 64 channel head/neck coil with all elements enabled. We collected 3 T1-weighted and 1 T2-weighted anatomical images using the Siemens product MPRAGE and Turbo Spin-Echo sequences (0.8 mm isotropic voxels, 256 X 240 mm slice FOV). The TE/TR of each anatomical image were: 2.24/2400 ms (T1w) and 564/3200 ms (T2w). These scans were used to coregister all functional images to each participant's native anatomic space. We acquired functional images using a Multiband (MB) 2D GE-EPI scanning sequence with MB factor of 4, acquiring 44 2.5 mm interleaved slides with no gap, with isotropic voxel size being 2.5 mm. The TE/TR of these scans were: 30/750 ms (4x simultaneous-multi-slice acceleration). We used freesurfer's recon-all script (version 6.0) to process all anatomical scans and AFNI's afni_proc.py command to preprocess all functional scans. Preprocessing included unwarping, alignment, rigid-body motion correction (6-parameter affine transform), and smoothing along the cortical surface normal. Retinotopic mapping data was spatially smoothed along the cortical surface (5 mm FWHM) for identifying ROIs, but feature-selective attention and spatial mapping task data was not smoothed beyond projecting to surface space and returning to volume space.

Region of interest definition

We identified 15 ROIs using independent retinotopic mapping data. We fit a vRF model for each voxel in the cortical surface (in volume space) using averaged and spatially smoothed (on the cortical surface; 5 mm FWHM) time series data across all retinotopy runs (8-12 per participant). We used a compressive spatial summation isotropic Gaussian model (Kay et al., 2013; Mackey et al., 2017) as implemented in a customized, GPU-optimized version of mrVista (see Mackey et al., 2017) for detailed description of the model). High-resolution stimulus masks were created (270 x 270 pixels) to ensure similar predicted responses within each bar size across all visual field positions. Model fitting began with an initial high-density grid search, followed by subsequent nonlinear optimization. Retinotopic ROIs (V1, V2, V3, V3AB, hV4, LO1, LO2, VO1, VO2, TO1, TO2, IPS0-3) were then delineated by projecting vRF best-fit polar angle and eccentricity parameters with variance explained ≥5% onto each participant's inflated cortical surfaces via AFNI and SUMA. We drew ROIs on each hemisphere's cortical surface based on previously-established polar angle reversal and foveal representation criteria (Amano et al., 2009; Mackey et al., 2017; Swisher et al., 2007; Wandell et al., 2007; Winawer & Witthoft, 2015). Finally, ROIs were projected back into

volume space to select voxels for analysis. For primary analyses (Figs. 4, 6), we focused on previously identified 'feature dimension maps' (Thayer & Sprague, 2023) which were aggregate ROIs made up of a combination of all voxels across retinotopic maps with similar feature selectivity (color: hV4, VO1, and VO2; motion: TO1 and TO2). For completeness, we additionally report results from each individual ROI (Figs. 5, 7).

Inverted encoding model.

We used a spatial inverted encoding model (IEM) to reconstruct images of stimulus-related activation patterns measured across entire ROIs (Sprague & Serences, 2013; Fig. 3). To do this, we first estimated an encoding model, which describes the sensitivity profile over the relevant feature dimension for each voxel in a region. This requires using data set aside for this purpose, referred to as the "training set". Here, we used data from the spatial mapping task as the independent training set. The encoding model across all voxels within a given region is then inverted to estimate a mapping used to transform novel activation patterns from a "test set" (runs from the feature-selective attention task) and reconstruct the spatial representation of the stimulus at each timepoint. For this analysis, we used voxels for which the best-fit vRF model explained ≥5% of the variance during the retinotopic mapping task.

We built an encoding model for spatial position based on a linear combination of 37 spatial filters (Sprague, Itthipuripat, et al., 2018; Thayer & Sprague, 2023). Each voxel's response was modeled as a weighted sum of each identically shaped spatial filter arrayed in a triangular grid (Fig. 3A). The centers of each filter were spaced by 2.83° and were Cosine functions raised to the 7th power:

$$f(r) = \left(0.5 + 0.5\cos\frac{\pi r}{s}\right)^7 \quad (1)$$

for $r < s$; 0 otherwise

where r is the distance from the filter center and s is a size constant. The size constant reflects the distance from the center of each spatial filter at which the filter returns to 0. This triangular grid of filters forms the basis set, or information channels for our analysis. For each stimulus used in our mapping task, we converted from a contrast mask to a set of filter activation levels by taking the dot product of the vectorized stimulus mask (n trials × p pixels) and the sensitivity profile of

each filter (*p* pixels × *k* channels). We then normalized the estimated filter activation levels such that the maximum activation was 1 and used this output as $\boldsymbol{C_1}$ in the following equation, which acts as the forward model of our measured fMRI signals:

$$B_1 = C_1 W \tag{2}$$

$\boldsymbol{B_1}$ (*n* trials × *m* voxels) in this equation is the measured fMRI activity of each voxel during the visuospatial mapping task and $\boldsymbol{W}$ is a weight matrix (*k* channels × *m* voxels) which quantifies the contribution of each information channel to each voxel. $\widehat{\boldsymbol{W}}$ can be estimated using ordinary least-squares linear regression to find the weights that minimize the differences between predicted values of $\boldsymbol{B}$ and the observed $\boldsymbol{B_1}$:

$$\widehat{W} = (C_1^T C_1)^{-1} C_1^T B_1 \tag{3}$$

This is computed for each voxel within a region independently, making this step univariate. The resulting $\widehat{W}$ represents all estimated voxel sensitivity profiles. We then used $\widehat{W}$ and the measured fMRI activity of each voxel (i.e., BOLD response) during each trial (using each TR from each trial, in turn) of the feature-selective attention task using the following equation:

$$\hat{C}_2 = B_2 \widehat{W}^T (\widehat{W} \widehat{W}^T)^{-1} \tag{4}$$

Here, $\hat{C}_2$ represents the estimated activation of each information channel (*n* trials × *k* channels) which gave rise to that observed activation pattern across all voxels within a given ROI ($\boldsymbol{B_2}$; *n* trials × *m* voxels). To aid with visualization, quantification, and coregistration of trials across stimulus positions, we computed spatial reconstructions using the output of Equation 4. To do this, we weighted each filter's spatial profile by the corresponding channel's reconstructed activation level and then sum all weighted filters together (Fig. 3B).

Since stimuli in the feature-selective attention task were randomly positioned on every trial, we rotated the center position of spatial filters such that the resulting 2D reconstructions of the stimuli were aligned across trials and participants (Fig. 3C). We then sorted trials based on condition (attend fixation, attend color, attend motion). Finally, we averaged the 2D reconstructions across trials within the same condition for individual participants, then across all participants for our grand-average spatial reconstructions (Fig. 6; Fig. 7). Individual values within

the 2D reconstructed spatial maps correspond to map activation at the corresponding relative visual field coordinates based on activation profiles across entire ROIs.

Critically, because we reconstructed all trials from all conditions of the feature-selective attention task using an identical spatial encoding model estimated with an independent spatial mapping task, we can compare reconstructions across conditions on the same footing (Sprague, Adam, et al., 2018; Sprague et al., 2019; Sprague, Itthipuripat, et al., 2018).

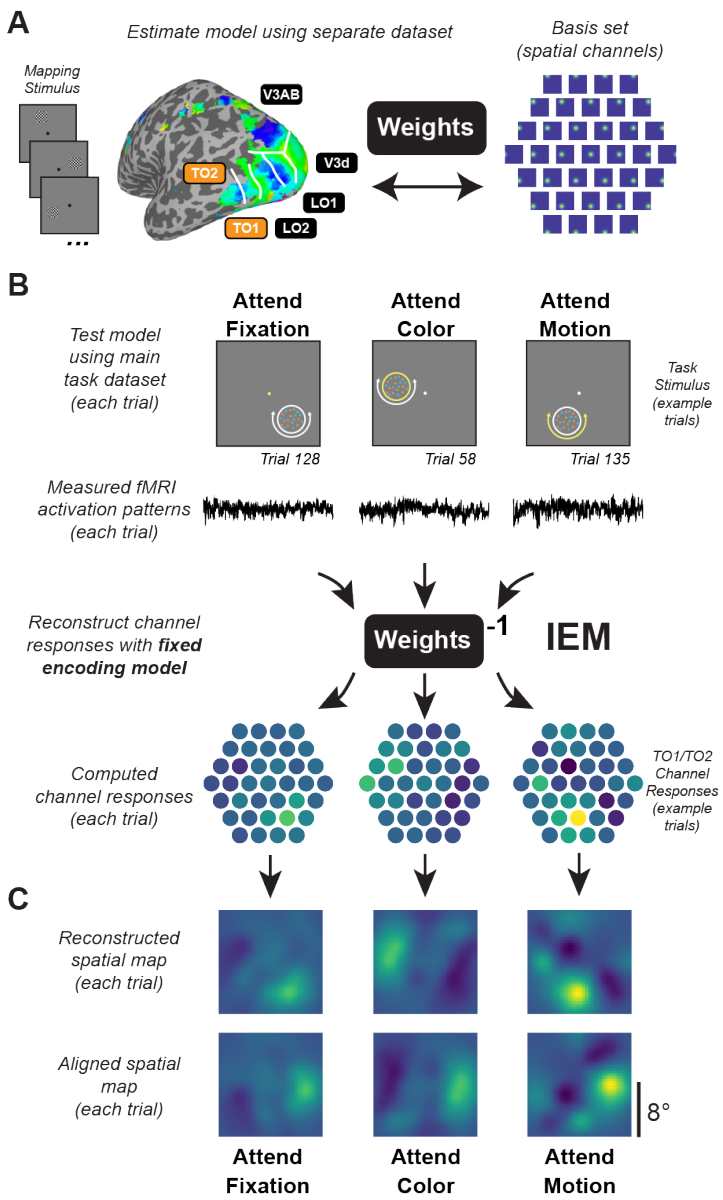

Figure 3: Inverted encoding model used to reconstruct maps of the visual field. A. Data from an independent spatial mapping task was used to estimate a spatial IEM for each ROI (for details, see Materials and Methods). **B.** Using this approach, activation patterns across all voxels in a region for each trial can be mapped onto the activation

of spatial channels that are then summed to produce reconstructed spatial maps of the visual field. Identical encoding models, estimated with the task in (A), are used to compute reconstructions for each condition, ensuring results can be compared on equal footing. We reconstructed maps for each trial using activation patterns averaged 4.5-7.5 s after cue onset and directly compared the map activation across conditions. **C.** Since maps are in visual field coordinates, we were able to rotate and align reconstructed spatial maps to the known position of the moving dot array, then average across trials within each condition. Computed channel responses are example data from a representative subject's (sub015) aggregate TO1/TO2 ROI.

Quantifying neural responses

First, to establish the impact of feature-selective attention to specific stimulus dimensions on overall activation in each ROI at the broadest spatial scale, we computed the average activation across all visually responsive voxels (defined as those with at least 5% variance explained based on vRF model fits) in each ROI (McMains et al., 2007; Runeson et al., 2013). Second, because we can measure activation across each entire region simultaneously (that is, voxels with the full spread of spatial encoding properties), we performed a multivariate spatial IEM analysis (described above) to transform activation patterns to reconstructed spatial maps for each trial. This analysis leverages all spatial information from all voxels simultaneously to recover the stimulus representation supported by the entire neural population distributed across each retinotopic region or set of regions (Fig. 3C). To quantify stimulus representations in reconstructed spatial maps, we computed the mean map activation of pixels located within a 1.5° radius disk centered at the known position of each stimulus (matching the stimulus radius of 1.5°; see Sprague, Itthipuripat, et al., 2018; Thayer & Sprague, 2023). This provides a single value corresponding to the activation of the stimulus location for a given condition, within each retinotopic ROI. To assess the spatial selectivity of reconstructed spatial maps, we compared the mean map activation at the location of stimulus to map activation at the location opposite fixation using a 1.5° radius disk (Figs. 6 and 7).

For univariate and multivariate analyses, we computed an attention modulation index (AMI) to directly compare the impact of attending to color vs motion on stimulus representations, while accounting for the overall impact of spatial attention, similar to previous reports (McMains et al., 2007). AMI was computed as the difference in activation between the attend motion and attend color conditions, where positive values indicate a stronger response to the stimulus when attending motion (we did not normalize by the sum because, in some regions/analyses, the color-

or motion-related activation was negative, making a normalized difference challenging to interpret).

In addition to the primary focused analyses on *a priori* feature-selective retinotopic regions (Fig. 4, 6), we also present all results computed for each retinotopic ROI individually (Figs. 5, 7).

Statistical analysis

Parametric statistical tests were conducted for all comparisons (repeated-measures ANOVAs and *t* tests). To account for possible non-normalities in our data, we generated null distributions for each test using a permutation procedure (see below) to derive p-values.

First, we used a one-way repeated-measures ANOVA (factor: attention condition, with 3 levels: attend fixation, attend color, attend motion) to determine whether behavioral accuracy on the fixation task depended on what was attended on each trial (Fig. 2B). We then conducted follow-up paired-samples *t* tests between all pairs of attention conditions to see which conditions differed. A paired-samples *t* test was conducted to compare attend color and attend motion RTs. The attend fixation condition was excluded for this analysis due to differences in stimulus onset (fixation target appeared briefly at least 500 ms after the dot array stimulus onset, while coherent color/motion was present throughout the dot stimulus presentation), making comparisons difficult to interpret.

When analyzing fMRI data, we focused on retinotopic regions that have previously been shown to be color- or motion-selective (Brewer et al., 2005; Huk et al., 2002; Thayer & Sprague, 2023). For univariate analyses, we computed two-way ANOVAs with attention condition (3 levels: attend fixation, color, and motion) and ROI (2 levels: color-selective hV4/VO1/VO2 and motion-selective TO1/TO2) as factors. Follow-up paired-samples *t* tests were then conducted to determine if responses differed within each ROI as a function of attention condition (Fig. 4, 5).

For multivariate neural responses, we first computed a three-way ANOVA with location (2 levels: stimulus location and opposite location), attention condition (3 levels: attend fixation, color, and motion), and ROI (2 levels: color-selective hV4/VO1/VO2 and motion-selective TO1/TO2) as factors. We then computed separate two-way ANOVAs for each relevant location (stimulus and opposite location), with attention condition (3 levels: attend fixation, color, motion) and ROI (2

levels: color-selective hV4/VO1/VO2 and motion-selective TO1/TO2) as factors. This analysis determined how task demands and feature selectivity uniquely modulated responses between regions at each location. Next, we computed the difference between stimulus and opposite locations and conducted a two-way ANOVA to determine the *spatial selectivity* of neural modulations with attention condition (3 levels: attend fixation, color, motion) and ROI (2 levels: color-selective hV4/VO1/VO2 and motion-selective TO1/TO2) as factors. To directly test whether feature-selective ROIs represent stimulus locations more strongly when goals direct attention to their preferred feature dimension, we computed a paired-samples *t* test on the AMI computed for color-selective and motion-selective ROIs (Fig. 6C, 7C). Lastly, one-sample *t* tests for univariate and multivariate AMIs were conducted for each region to determine if neural responses were feature selective (AMI different from zero). When performing any ANOVA across individual ROIs with ROI as a factor, there were 13 levels (all other ROIs: see 'region of interest definition' above).

For our shuffling procedure, we used a random number generator that was seeded with a single value for all analyses. The seed number was randomly selected using an online random number generator (https://numbergenerator.org/random-8-digit-number-generator). Within each participant, averaged data within each condition were shuffled across conditions for each participant individually, and once shuffled, the statistical test of interest was recomputed over 1000 iterations. *P* values were derived by computing the percentage of shuffled test statistics that were greater than or equal to the measured test statistic. We controlled for multiple comparisons using the false discovery rate (Benjamini & Yekutieli, 2001) across all comparisons within an analysis when necessary. Error bars are standard error, unless noted otherwise.

Data & Code Availability

All data supporting the conclusions of this report and all associated analysis scripts are available on Open Science Framework (https://osf.io/3qrpc/?view_only=93818b413458400f9fac93dd878f09a2). To protect participant privacy, and in accordance with IRB-approved procedures, freely available data is limited to extracted timeseries for each voxel of each ROI for all scans of the study. Whole-brain 'raw' data will be made available from the authors upon reasonable request.

Results

Behavior

Participants performed a feature-selective attention task while we used fMRI to measure BOLD activation patterns across several independently defined retinotopic regions. We cued participants to covertly attend to either the color or motion of a colored-moving dot stimulus on every trial to determine which individual feature value was most prevalent. Accuracy differed across attention conditions as indicated by a permuted one-way repeated measures ANOVA (p = 0.002), where accuracy was worse on attend color than attend motion trials (paired sample t-test, $p < 0.001$). There was no significant difference in response times (RT) between the attend color and motion conditions (p = 0.270).

Univariate neural responses

First, to determine whether the attended feature dimension impacts aggregate neural responses across independently mapped retinotopic ROIs selective for color and motion, we compared univariate neural responses averaged across all visually responsive voxels within each region, without considering their spatial selectivity with respect to the stimulus location. Hemodynamic response functions (HRFs) for each aggregate feature-selective ROI show a stimulus-related visually evoked response which differs as a function of task condition (Fig. 4A). To quantify these attentional modulations, we averaged responses 4.5-7.5 seconds after the attention cue was presented (3-6 s after stimulus onset; Fig. 4 B-C). A two-way permuted ANOVA with attention condition (attend fixation, color, and motion), and ROI (hV4/VO1/VO2 and TO1/TO2) as factors revealed a main effect of ROI ($p < 0.001$) as well as a two-way interaction ($p < 0.001$). To determine if each region's overall activation was modulated by attention to the peripheral stimulus as compared to attention to the fixation point, we performed follow-up permuted two-sample t tests within each ROI between each attend-stimulus condition and the attend fixation condition. This analysis provides insights into how directing attention to the location of the peripheral stimulus may modulate overall responses across all voxels within each region, possibly by enhancing the response to goal-relevant feature dimensions. In hV4/VO1/VO2, there was a stronger response when attending to the stimulus than when attending to fixation (fixation vs color: $p < 0.001$; fixation vs motion: p = 0.001), indicating that attention directed to the stimulus enhanced

responses regardless of attended feature. However, TO1/TO2 showed no general modulation when attention was directed to the stimulus for either condition when compared to attending fixation (fixation vs color: $p = 0.067$; fixation vs motion: $p = 0.598$).

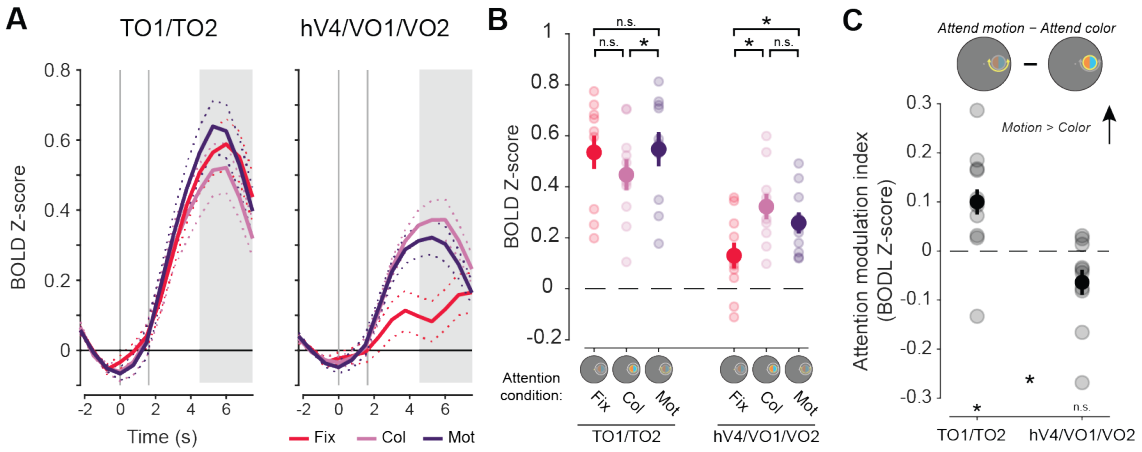

Figure 4: BOLD response across all voxels in aggregate feature-selective visual cortex ROIs show motion selectivity for TO1/TO2 and task demand modulation in hV4/VO1/VO2. A. HRFs from aggregate feature-selective retinotopic regions TO1/TO2 and hV4/VO1/VO2. Gray lines indicate cue (first) and stimulus (second) onsets. Shaded area indicates time points used for trial averaging. Dot array stimulus was presented 1.5 to 4.5 seconds (second gray line to start of shaded area). **B.** Mean BOLD response for each attention condition. Motion-selective regions showed no modulation of response when attending to either feature of the stimulus (color or motion) relative to when attending to fixation (color vs fixation: $p = 0.07$; motion vs fixation: $p = 0.59$). Color-selective regions were modulated by the task-relevance of the peripheral stimulus, with a greater region-level response when attending to the color ($p < 0.01$) or the motion ($p < 0.01$) relative to attending to fixation. **C.** Attention modulation index for both aggregate feature-selective regions. Only motion-selective regions had a response that differed from zero ($p = 0.02$). * indicate significant differences based on permuted paired or single sample *t* tests ($p < 0.05$). Error bars indicate SEM within participants. Col: color, Mot: motion, Fix: fixation.

Next, to assess whether the specific attended feature modulates activation in each region, we computed an attention modulation index (AMI) by subtracting the neural responses between the attend motion and attend color conditions, with a positive value indicating a stronger response when attending to motion (note that this is equivalent to a paired comparison between attend-motion and attend-color conditions). Permuted one-sample *t* tests revealed that TO1/TO2 had a significant positive value ($p = 0.024$). Even though motion-selective regions showed no general modulation when attention was directed to the stimulus as compared to fixation (Fig. 4B), there was a difference in overall response as a function of attended feature dimension. In contrast, hV4/VO1/VO2 AMI based on overall activation did not significantly differ from zero ($p = 0.061$). Thus, while color regions were more active when the peripheral stimulus was relevant, their

overall activation did not significantly vary based on what feature dimension was attended. Finally, a direct comparison of AMI between color- and motion-selective regions revealed a significant difference (Fig. 4C; $p < 0.001$), consistent with a region-selective modulation of overall activation levels based on attended feature.

As an additional exploratory analysis to assess the specificity of our results across all retinotopic regions identified, we repeated these analyses for each ROI individually (for all statistics, see Table 1). Results were generally consistent with the aggregate ROIs. First, the two-way permuted ANOVA with factors of ROI and attention condition revealed a main effect of ROI ($p < 0.001$) and an interaction ($p = 0.013$). For each color-selective region, follow-up permuted two-sample *t* tests showed modulations based on stimulus relevance: hV4 (fixation vs color: $p = 0.002$; fixation vs motion: $p = 0.013$); VO1 (fixation vs color: $p < 0.001$; fixation vs motion: $p < 0.001$); and VO2 (fixation vs color: $p < 0.001$; fixation vs motion: $p = 0.004$). TO1 (fixation vs color: $p = 0.090$; fixation vs motion: $p = 0.234$) and TO2 (fixation vs color: $p = 0.035$; fixation vs motion: $p = 0.771$) were not modulated by stimulus relevance. Permuted one-sample *t* tests indicate that AMIs for TO1 ($p = 0.024$) and TO2 ($p = 0.014$) had trending (defined as $p < 0.05$, uncorrected) negative values but these did not survive FDR correction across 13 ROIs. Similar to the aggregate ROI analyses, the AMI value from hV4 was not different from zero ($p = 0.224$). However, both VO1 ($p = 0.022$) and VO2 ($p = 0.018$) had trending negative values (failed FDR corrections), providing evidence that these regions were more strongly modulated by attending to stimulus color. These results are consistent with a previous study by McMains and colleagues (2007), which also found that TEO (corresponding to VO1/2; Kravitz et al, 2013) was the only color-selective region with increased responses when color was attended, perhaps indicating a specialized role of higher-order, more anterior ventral regions in modulating processing based on attention. Additionally, looking at individual motion-selective retinotopic regions, we see similar modulations as the aggregate ROI, with greater responses when attending to motion than color. It is notable that only color and motion regions had any evidence for feature selectivity, as all other retinotopic regions AMI were no different from 0 (all *p*'s > 0.05; see Table 1).

In sum, attending to specific feature dimensions rescales the overall activation levels of feature-selective retinotopic ROIs, with regions showing greater univariate activation when their preferred

feature dimension is task-relevant. However, this univariate analysis disregards the spatially selective nature of neural codes in retinotopic ROIs. Next, we asked: is this re-scaling driven primarily by neural populations encoding the attended location, or instead, does this reflect a global scaling of activation across the entire ROI (Fig. 1)?

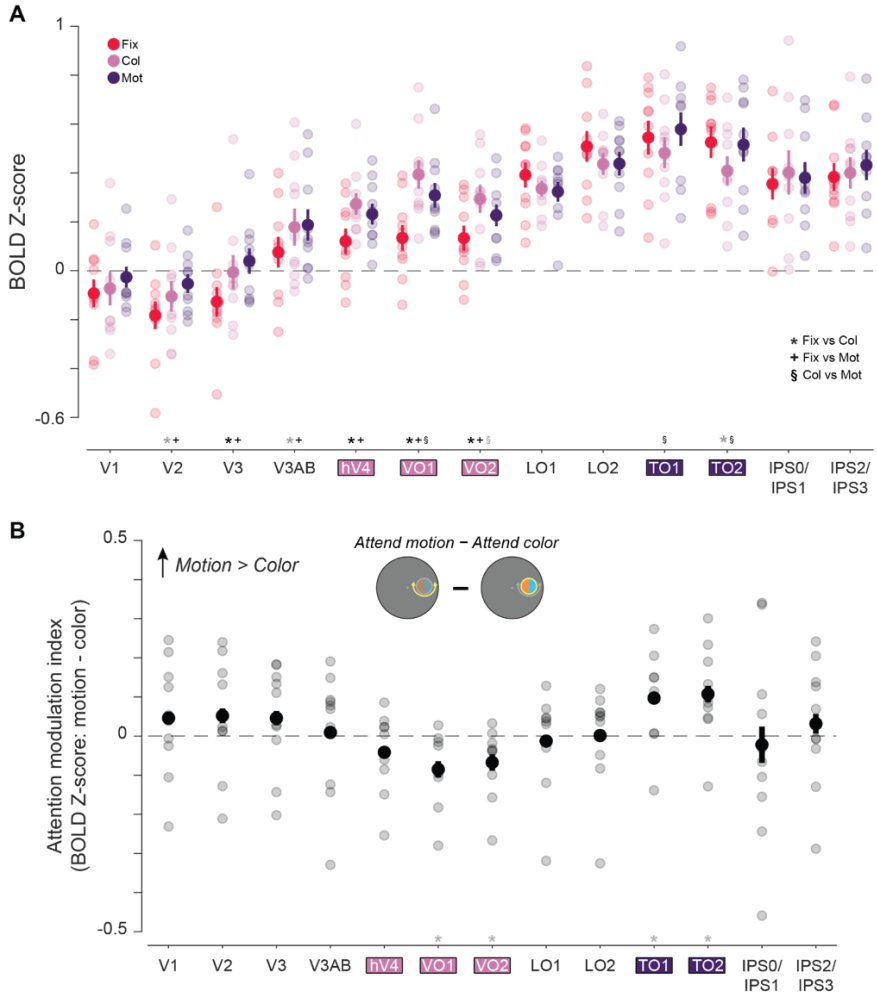

Figure 5: BOLD response across all voxels for individual regions of interest show selective attentional modulation based on feature preferences. A. Mean BOLD response for each attention condition from all individual retinotopic regions (feature-selective retinotopic regions hV4, VO1, VO2, TO1, and TO2, presented in aggregate in Fig. 4, are highlighted). Both individual motion-selective regions showed no response modulation when attending to either feature of the stimulus (color or motion) relative to when attending to fixation (all *p*'s > 0.05). Each color-selective region was modulated by task demands, with a greater region-level response when attending to the color ($p < 0.05$) or the motion ($p < 0.05$) relative to attending to fixation. **B.** Attention modulation index for each individual feature-selective region. Both motion-selective regions had a positive response that differed from zero (TO1 and TO2: $p < 0.05$, uncorrected), indicating a stronger response when attending to motion. Color-selective regions VO1 and VO2 had a negative response that differed from zero (VO1 and VO2: $p < 0.05$, uncorrected), indicating a preference for color. However, hV4 did not differ from zero ($p > 0.05$). * indicate significant differences based on permuted paired or single sample *t* tests ($p < 0.05$), corrected for multiple comparisons via FDR across 13 ROIs. Gray asterisks indicate trends, defined based on p < 0.05 without correction. Error bars indicate SEM within participants. Col: color, Mot: motion, Fix: fixation.

Figures 4, 5	Attend Fixation	Attend Color	Attend Motion	Fixation vs Color	Fixation vs Motion	Color vs Motion	AMI
V1	0.122	0.333	0.582	0.605	0.151	0.343	0.365
V2	**0.006**	0.138	0.176	*0.042*	**0.005**	0.308	0.286
V3	0.061	0.938	0.483	**0.006**	**0**	0.321	0.327
V3ab	0.249	0.041	**0.012**	*0.032*	**0.018**	0.848	0.874
hV4	0.054	**0**	**0**	**0.002**	**0.013**	0.212	0.224
VO1	**0.023**	**0**	**0**	**0**	**0**	**0.016**	*0.013*
VO2	*0.027*	**0**	**0**	**0**	**0.004**	*0.031*	*0.029*
LO1	**0**	**0**	**0**	0.234	0.073	0.781	0.792
LO2	**0**	**0**	**0**	0.204	0.106	0.980	0.974
TO1	**0**	**0**	**0**	0.090	0.234	**0.020**	*0.024*
TO2	**0**	**0**	**0**	*0.035*	0.771	**0.017**	*0.014*
TO1/TO2	**0**	**0**	**0**	0.067	0.598	**0.014**	**0.024**
hV4/VO1/VO2	**0.029**	**0**	**0**	**0**	**0**	0.055	0.061
IPS0/IPS1	**0**	**0**	**0**	0.506	0.617	0.783	0.788
IPS2/IPS3	**0**	**0**	**0**	0.758	0.302	0.536	0.555

Table 1. Univariate statistics. *P* values for univariate comparisons for individual attention conditions and between specific attention conditions for each ROI. Comparisons include paired-sample *t* tests between the attend fixation, color, and motion conditions. One-sample *t* tests for individual attention conditions and AMI are also included. Even though the Color vs Motion and AMI analysis are effectively the same comparison, *p* values are slightly different due to the random shuffling procedure used to generate non-parametric distributions for statistical comparison (see statistical procedures above) and different number of comparisons when doing FDR corrections (78 for Color/Motion/Fixation comparisons; 13 for AMI). Bold denotes significance after FDR correction for all comparisons (q = 0.05). Significance before FDR corrections using α = 0.05 denoted with italics.

Multivariate spatial representations

While informative, the univariate results alone cannot speak to the mechanism of attentional modulation of feature dimension maps, because univariate analyses collapsing across entire retinotopic maps are unable to discern which parts of the region—corresponding to different locations in the visual field—are modulated by task demands. To further evaluate how feature dimension maps encoded within TO1/TO2 and hV4/VO1/VO2 are modulated by task demands, we computed an IEM for each region (Fig. 3; see Materials and Methods; Sprague et al., 2018). To do this, we used an independent "mapping" task to estimate a spatial encoding model for each voxel. The model was parameterized as a series of weights across several smooth, overlapping spatial channels (Fig. 3A). The weights from the spatial encoding model across all voxels that comprised a feature-selective ROI were then inverted which allowed us to reconstruct spatial maps based on activation profiles from the feature-selective attention task (Fig. 3B-C) averaged over 4.5-7.5 s after cue onset.

The reconstructed spatial maps (Fig. 6A) appear to have a stronger response at the aligned stimulus location than the rest of the modelled visual field (for individual ROI results, see Fig. 7). We quantified whether there was a stronger response to the stimulus by averaging map activation profiles within a 1.5° radius disk (stimulus size) at the aligned stimulus location. We could similarly assess map activation at the opposite screen location across fixation, which had no stimulus present (and was never attended; Fig. 6B). A three-way permuted ANOVA with location (stimulus and opposite), attention condition (attend fixation, color, and motion), and ROI (hV4/VO1/VO2 and TO1/TO2) revealed a three-way interaction ($p < 0.001$). We conducted two-way permuted ANOVAs for each location with ROI (hV4/VO1/VO2 and TO1/TO2) and attention condition (attend fixation, color, and motion) as factors to better determine what was driving the three-way interaction. Like the univariate approach, the aim of this analysis was to determine whether motion- and color-selective regions were modulated by task demands and to determine if these regions differed in their representation depending on which feature dimension was prioritized. Responses at the stimulus location showed a main effect of attention condition ($p < 0.001$) and ROI ($p < 0.025$) as well as an interaction ($p < 0.001$). Follow-up permuted paired sample *t* tests indicate that color-selective regions hV4/VO1/VO2 have a stronger neural response when attending to stimulus color than when attending to fixation ($p < 0.001$). There was also a stronger response when attending to stimulus motion than fixation ($p < 0.001$), consistent with spatial attention allocated to the stimulus location. Similarly, motion-selective regions TO1/TO2 responded more strongly to the stimulus when attending to either feature than when attending to fixation (fixation vs color: $p < 0.001$; fixation vs motion: $p < 0.001$), again due to spatial attention.

A two-way ANOVA at the opposite location had no effect of ROI ($p = 0.131$) but did have a main effect of attention condition ($p < 0.001$) and an interaction between ROI and attention condition ($p = 0.022$). Follow-up permuted two-sample *t* tests reveal that there was a greater response at the opposite location in motion-selective regions when attending to fixation than either feature dimension of the stimulus (fixation vs color: $p = 0.001$; fixation vs motion: $p = 0.001$). Color-selective regions had a greater response at the opposite location when attending to fixation than attending to motion ($p = 0.008$). Thus, particularly for motion-selective regions, there are some

task-dependent modulations occurring at the location opposite of the stimulus in reconstructed maps.

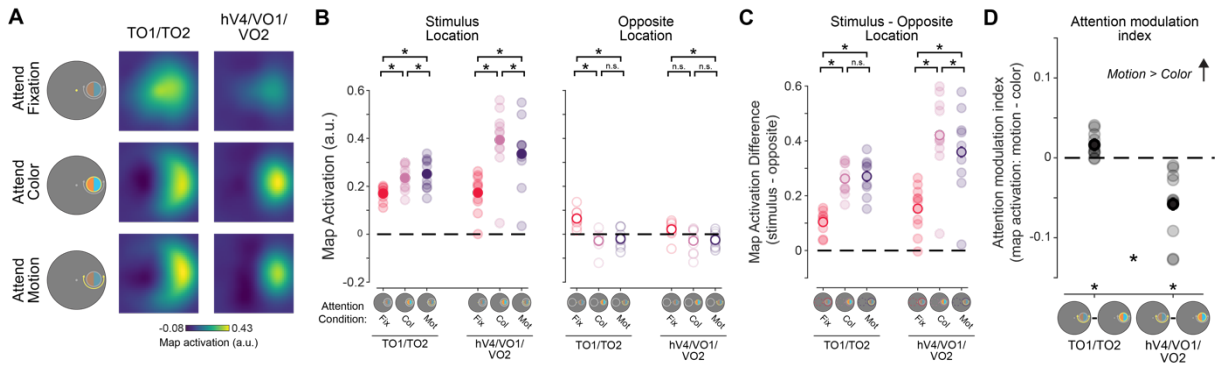

Figure 6: Reconstructed spatial maps from color- and motion-selective regions are differentially modulated by feature-selective attention. A. Reconstructions from motion- and color-selective regions from the attend fixation, color, and motion conditions. Qualitatively, there is a strong response at the aligned location, particularly when attending to either the motion or color of the stimulus. **B.** We determined whether each region was modulated by task demands by assessing the strength of the stimulus representation for each stimulus condition as well as the opposite location. For both regions, the neural response at the stimulus location was stronger when attending to either feature of the stimulus as compared to attend fixation trials, indicating that stimulus representations in these regions track task-relevance of the stimulus location. At the opposite location, responses were stronger when attending to fixation. **C.** Difference in map activation between the stimulus and opposite locations for each stimulus condition, which indexes spatial specificity of map activation profiles. Responses were always greater when attending to the stimulus as compared to fixation for both regions. While color regions hV4/VO1/VO2 show a difference between the attend motion and color conditions (p= 0.02), indicating local enhancement, motion regions TO1/TO2 showed no difference (p = 0.61), indicating global enhancement. **D.** Difference in map activation at the stimulus location between the attend motion and attend color conditions. Responses reliably differed from zero for both ROIs, where motion-selective areas had stronger response when attending to motion and color regions had a stronger response when attending to color. Responses differed between regions. * Significant difference based on permuted paired-samples t test ($p < 0.05$). Error bars indicate SEM within participants.

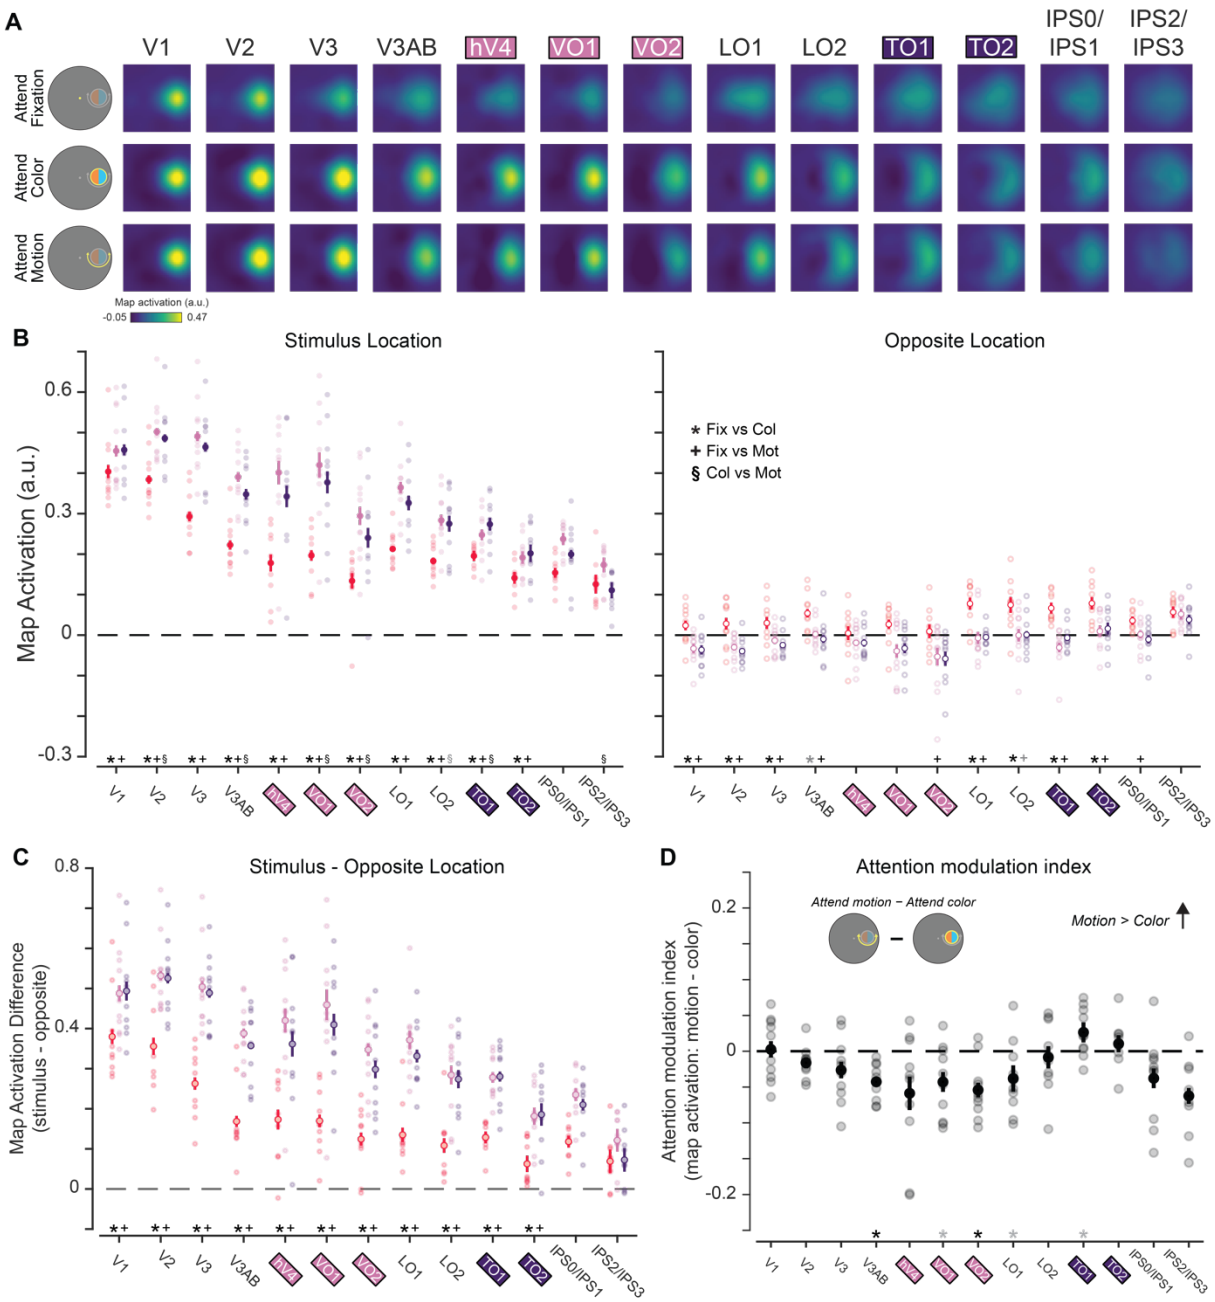

Figure 7: Reconstructed spatial maps for individual retinotopic ROIs. A. Reconstructions from all individual retinotopic ROIs, for the attend fixation, color, and motion conditions. Qualitatively, there is a strong response at the aligned location, particularly when attending to either the motion or color of the stimulus. **B.** We determined whether each region was modulated by task demands by assessing the strength of the stimulus representation for each stimulus condition as well as the opposite location. In all feature-selective regions, the neural response at the stimulus location was stronger when attending to either feature of the stimulus as compared to attend fixation trials, indicating that each region tracks task-relevance of the stimulus location. At the opposite location, responses were often strongest when attending to fixation. **C.** Difference in map activation between the stimulus and opposite locations for each stimulus condition. Responses were always greater when attending to either the color or motion of the stimulus for all regions (all p's < 0.05). **D.** Difference in map activation between the attend motion and attend color conditions. Responses reliably differed from zero for TO1, VO1, and VO2, where TO1 had stronger response when attending to motion and VO1 and VO2 had stronger responses when attending to color. * Significant

difference based on permuted paired-samples *t* test ($p < 0.05$), corrected for multiple comparisons via FDR. Gray asterisks indicate trends, defined as $p < 0.05$, uncorrected. Error bars indicate SEM within participants.

	Stimulus Location			**Opposite Location**			**Location Difference**			
Figures 6, 7	Fixation vs Color	Fixation vs Motion	Color vs Motion	Fixation vs Color	Fixation vs Motion	Color vs Motion	Fixation vs Color	Fixation vs Motion	Color vs Motion	AMI
V1	**0.023**	**0.004**	0.856	**0**	**0**	0.79	**0.003**	**0**	0.778	0.853
V2	**0**	**0**	**0.036**	**0.011**	**0.003**	0.294	**0**	**0**	0.619	0.055
V3	**0**	**0**	0.114	**0.017**	**0**	0.241	**0**	**0**	0.324	0.11
V3ab	**0**	**0**	**0**	*0.044*	**0.006**	0.616	**0**	**0**	0.185	**0**
hV4	**0**	**0**	0.055	0.351	0.234	0.991	**0**	**0**	0.087	0.064
VO1	**0**	**0**	**0.023**	0.084	0.057	0.701	**0**	**0**	0.099	*0.024*
VO2	**0**	**0.006**	**0.002**	0.066	**0.003**	0.83	**0**	**0.002**	0.106	**0**
LO1	**0**	**0**	*0.048*	**0.002**	**0.003**	0.919	**0**	**0**	0.157	*0.047*
LO2	**0**	**0.008**	0.588	**0.015**	*0.037*	0.949	**0**	**0.007**	0.613	0.623
TO1	**0.007**	**0**	**0.042**	**0.005**	**0.002**	0.240	**0**	**0**	0.912	*0.024*
TO2	**0**	**0.005**	0.312	**0.011**	**0.008**	0.749	**0**	**0**	0.881	0.353
TO1/TO2	**0**	**0**	**0.011**	**0**	**0**	0.619	**0**	**0**	0.610	**0.006**
hV4/VO1/VO2	**0**	**0**	**0.005**	0.040	**0.008**	0.748	**0**	**0**	**0.019**	**0**
IPS0/IPS1	**0**	0.071	0.080	0.084	**0**	0.555	**0**	**0.005**	0.280	0.076
IPS2/IPS3	0.085	0.623	**0.004**	0.523	0.277	0.422	0.084	0.908	0.055	**0.005**

Table 2. Multivariate Statistics. *P* values for IEM comparisons between specific attention conditions at certain locations in the visual field for each ROI. Comparisons include paired-sample *t* tests between the attend fixation, color, and motion conditions. One-sample *t* tests for AMI are also included. Even though the Color vs Motion at the stimulus location and AMI analyses are effectively the same comparison, *p* values are slightly different due to the random shuffling procedure used to generate non-parametric distributions for statistical comparison (see statistical procedures above) and different number of comparisons when doing FDR corrections (39 for Fixation/Color/Motion comparisons; 13 for AMI). Bold denotes significance after FDR correction for all comparisons ($q = 0.05$). Significance before FDR corrections using $\alpha = 0.05$ denoted with italics.

Next, to evaluate the spatial selectivity of these modulations, we computed the difference in map activation between the stimulus and opposite locations for each attention condition (Fig. 6C). A 2-way repeated measures ANOVA with attention condition (attend fixation, color, and motion) and ROI (TO1/TO2 and hV4/VO1/VO2) as factors reveal a main effect of ROI ($p = 0.018$), attention condition ($p < 0.001$), and an interaction between ROI and attention condition ($p = 0.006$). Follow up comparisons indicate that in color-selective regions, the stimulus location was more strongly represented when attending to a stimulus feature (fixation vs color: $p < 0.001$; fixation vs motion: $p < 0.001$). Furthermore, the stimulus location was more strongly represented when attending to

color than when attending to motion (p = 0.019), indicating that feature-selective attention resulted in localized enhancement of the location of the stimulus representation. Motion regions also exhibit a stronger stimulus representation when attention was directed to the stimulus (attend color or attention motion > attend fixation, each $p < 0.001$). However, there was no difference between the attend color and motion conditions in the spatially selective stimulus representation ($p = 0.610$). Therefore, while TO1/TO2 is spatially selective, in the sense that spatial attention to the stimulus location locally enhances its representation, feature-selective attention results in a global shift in activation across the entire map (including the opposite location with no stimulus present). This striking and unexpected divergence between the modulatory profiles across color- and motion-selective regions suggests there may be unique spatial profiles for attentional modulation by feature relevance in color- and motion-selective regions.

Finally, we compared the AMI at the stimulus location (Fig. 6D), which indexes the difference in activation across feature-selective attention conditions, within and between color- and motion-selective ROIs. The motion-selective ROIs TO1/TO2 showed a significantly negative AMI (such that the stimulus representation was stronger during attend-motion than attend-color trials, p = 0.006), and the color-selective ROIs hV4/VO1/VO2 showed a significantly positive AMI ($p < 0.001$). Importantly, the AMIs derived from each set of regions significantly differed ($p < 0.001$). Altogether, this result supports a dissociable modulation of activation profiles across color- and motion-selective ROIs based on the attended feature dimension, with task-relevant feature-selective ROIs more strongly encoding stimulus locations when they are relevant for task performance.

As an exploratory analysis, we performed all multivariate IEM analyses on each individual retinotopic ROI (Fig. 7, Table 2). A three-way ANOVA with location (stimulus and opposite), attention condition (attend fixation, color, and motion), and ROI (hV4/VO1/VO2 and TO1/TO2) revealed a three-way interaction ($p < 0.001$). Follow-up two-way ANOVAs for each location (Fig. 7B) had main effects of condition and ROI (all p's < 0.001). At the stimulus location, the interaction was significant ($p < 0.001$), but at the opposite location there was no interaction ($p = 0.121$). A two-way ANOVA on the activation difference between locations (Fig. 7C) with attention condition (attend fixation, color, and motion) and ROI (hV4/VO1/VO2 and TO1/TO2) as factors revealed a

main effect of condition ($p < 0.001$) and ROI ($p < 0.001$) as well as an interaction ($p < 0.001$). Follow-up permuted paired sample *t* tests are generally consistent with aggregate feature-selective ROIs: neural responses were always greater when attending to either feature of the stimulus (color or motion) than when attending to fixation (all *p*'s < 0.05), demonstrating that the stimulus representation was always stronger when it was relevant to the ongoing task. However, responses in individual color-selective ROIs differed from the aggregate color ROI. Responses were no different between the attend color and attend motion conditions (hV4: $p = 0.087$; VO1: $p = 0.099$; VO2: $p = 0.106$), but all regions were qualitatively consistent with the aggregate result, in that color was numerically greater than motion. Thus, integrating responses across regions likely allowed the localized feature-selectivity to be identified for color regions. Only some individual feature-selective ROIs had AMIs that were trending or significantly differed from zero (TO1: $p = 0.024$, uncorrected; VO1: $p = 0.024$, uncorrected; VO2: $p = 0.001$). Notably, hV4 was not feature selective based on this analysis ($p = 0.064$), which is consistent with the univariate AMI results suggesting that this region may be weakly modulated by task demands. TO2 was also not different from zero ($p = 0.353$). In sum, analyzing individual retinotopic ROIs reveals a gradient of attentional modulations, with individual color-selective (hV4, VO1, VO2) and motion-selective (TO1, TO2) showing qualitatively extreme modulations when attending to color vs motion, and suggests that specific retinotopic subregions of our *a priori* defined aggregate ROIs (Fig. 6) contribute to our primary results.

Discussion

The aim of the current study was to establish the spatial profile of goal-driven feature-selective attention within neural feature dimension maps (Fig. 1). We presented a colorful moving dot stimulus on each trial, but cued participants to attend to either the color or motion of that stimulus (Fig. 2). Our design isolated the influence of observer goals because the goal-relevant feature dimension could change on each trial, but the visual input was consistent across the entire experiment. Region-level univariate methods (Fig. 4/5) showed unique modulations for color- and motion-selective regions: motion regions were feature-selective, but not modulated by attention directed to a peripheral stimulus as compared to fixation; color-selective regions were modulated

by covert spatial attention, but this modulation was not feature-selective. However, by using an inverted encoding model (Fig. 3) to reconstruct spatial maps from both motion- and color-selective regions, we were able to show that stimulus representations in both regions were modulated by task demands and were strongest when attending to each region's preferred feature (Fig. 6/7). Interestingly, different spatial profiles were observed for each feature-selective region: enhancement in color regions was *localized* to the stimulus, while motion regions were *globally* modulated, such that locations containing no stimulus were enhanced when motion was relevant. Not only do these findings support the role of each region as a 'neural feature dimension map' (Thayer & Sprague, 2023), they extend previous work by characterizing the unique spatial profiles that emerge through attention in color- and motion-selective regions.

Previous univariate fMRI studies average across voxels with variable spatial selectivity (Kriegeskorte & Bandettini, 2007; Serences & Saproo, 2012), making it difficult to ascertain the spatial profile of attentional modulations across visual field maps encoded by feature-selective regions. However, the multivariate approach implemented here revealed that color and motion regions have unique spatial prioritization profiles. Namely, neural responses in color regions primarily differed across attention conditions at the stimulus location (Fig. 6B-D), consistent with a local enhancement of stimulus processing at task-relevant locations (Fig. 1). Motion regions also showed a stronger response at the stimulus location when motion was attended (Fig. 6B/D), along with a similar bolstering of the response to the opposite, task-irrelevant location (Fig. 6B/C). This is consistent with a global enhancement of motion processing across all visual field locations when motion is made task-relevant (Fig. 1), similar to previous findings for attention to a specific motion value (Saenz et al., 2002; Serences & Boynton, 2007).

The present findings are likely due to feature-based attention (Maunsell & Treue, 2006; Scolari et al., 2014), which typically results in a spatially-global enhancement of neural responses associated with attended feature values. This enhancement occurs when a feature value is relevant, resulting in an increased response from all feature-selective neural populations, regardless of spatial selectivity, and drives a spatially global enhancement that has been observed for color (McAdams & Maunsell, 2000; Saenz et al., 2002; Zhang & Luck, 2009), motion (Martínez-Trujillo & Treue, 2002; Schneider, 2011; Schoenfeld et al., 2007; Serences & Boynton, 2007; Treue & Trujillo, 1999),

and orientation (Ester et al., 2016; McAdams & Maunsell, 2000). Notably, this finding is primarily observed when there are multiple stimuli at distinct locations, where one stimulus is relevant and another is irrelevant to current goals and is strongest at the irrelevant stimulus location when the feature value matches the relevant stimulus feature. If feature-based attention modulates responses across the visual field, then all neural populations tuned to the relevant feature should have increased firing, even when there is no stimulation. Serences & Boynton (2007) found evidence for successful decoding of the attended motion direction when motion was goal-relevant from voxels with no visual input. However, similar demonstrations that attended feature values can be decoded from stimulus-absent locations in the visual field are absent for attention to stimulus color, even though much of the evidence supporting feature-based attention uses colored stimuli. Our results may provide unique insight. The IEM approach employed here allowed for a demonstration of global modulation by feature-selective attention in motion-sensitive regions TO1/TO2, and so the absence of a similar observation for color-sensitive regions hV4/VO1/VO2 offers compelling evidence for a primarily local modulation of stimulus representations when color was task relevant. For these color-selective regions, it could be that global modulation is only observable under specific circumstances, such as when multiple stimuli are present, as is commonly the case in other studies requiring selective attention to a specific color value (e.g., Saenz et al., 2002).

Competing models make divergent predictions as to how task relevance of specific feature dimensions impacts activation profiles over feature dimension maps. Thus, it is necessary to characterize activation profiles across the entire visual field (Fig. 3) to finely assess properties of feature dimension maps and adjudicate between these models. 'Dimensional weighting' models claim relevance modulates responses across entire feature dimension maps in a spatially global manner by enhancing activation across the entire feature dimension map (Found & Müller, 1996; Liesefeld & Müller, 2019, 2023; Müller et al., 1995), which is consistent with our results from TO1/TO2. Other models claim that a spatial 'attentional window' restricts how much of the visual field is prioritized (Theeuwes, 2010, 2023), which is consistent with spatial activation profiles in hV4/VO1/VO2. As such, it appears aspects of both models are supported by the current findings. However, as whole, results from the current study have interesting parallels to the dimension

weighting priority map model. Dimensional weighting appears to occur more weakly when color is the relevant feature dimension (Liesefeld, Liesefeld, Pollmann, et al., 2019). It has been argued that this is because color is comprised of multiple dimensions (e.g., luminance, red-green, and blue-yellow in CIE *Lab* color space). Rather than considering color to uniquely contain multiple dimensions, as motion is similarly multi-dimensional (e.g., velocity, direction, and complex motion patterns), it may be that color is unique for other reasons, including peculiarities of neural architectures supporting local processing within individual feature dimension maps. Perhaps the local modulations we observed in hV4/VO1/VO2 are the neural substrate of the finding that task-driven enhancement of color acts on a more localized portion of the visual field, whereas our observation of more global enhancement in TO1/TO2 when motion is attended supports robust dimensional weighting across all spatial locations that is typically observed (Liesefeld, Liesefeld, & Müller, 2019; Liesefeld & Müller, 2019; Sauter et al., 2018).

One possibility as to why motion regions implement global modulations, while color regions locally enhance responses, is that vRFs in TO1/TO2 are much larger than hV4/VO1/VO2 (Wandell & Winawer, 2015). Thus, motion regions may necessarily enhance processing across large swaths of the visual field when motion is task relevant. This difference in encoding properties between regions may stem from the idiosyncrasies of color and motion as feature dimensions. By definition, to compute motion direction, the start and end point of an item are used to generate a direction vector (Albright & Stoner, 1995; Borst & Egelhaaf, 1989). Therefore, the global enhancement observed in the current study may be due to the specialization of TO1/TO2 in processing motion, a unique feature dimension which includes a temporal component. Notably, there was no requirement to prioritize a single location over others in the current study because only a single stimulus was ever presented. It is possible that different results would emerge when multiple stimuli are competing for attentional priority (Desimone & Duncan, 1995; Harrison et al., 2024; Sprague, Itthipuripat, et al., 2018).

Previous reports (McMains et al., 2007; Runeson et al., 2013) investigated how individual color- and motion-selective regions are modulated by task demands. Our findings were remarkably consistent with those found by McMains and colleagues (2007). Univariate analyses in color-selective region hV4 showed spatial selectivity, while the motion-selective region was not spatially

selective. Note, that the spatial analysis conducted by McMains et al. was very coarse, as it was only able to show that the hemisphere contralateral to the stimulus had a greater response than the ipsilateral hemisphere. To assay the full spatial profile of activation relative to the stimulus location, multivariate techniques like the one implemented in the current study must be used. Importantly, our IEM approach allowed for a fully randomized stimulus position on each trial, minimizing the impact of spatial expectation effects, and minimizing the impact of any idiosyncrasies of cortical sampling of different visual field locations. Additional similarities between our results and those of McMains et al (2007) are present as there was no significant modulation based on attended feature dimension in hV4, such that there was no greater univariate response when color was goal relevant compared to motion (Fig. 5). Also converging with their results, in our study, the more anterior color-selective regions VO1 and VO2 (parallel to TEO as described in their study; Brewer et al., 2005) showed the primary evidence for color-selective attentional modulations, with statistically less-significant effects observed in V4 (Figs. 5,7). Based on this finding, it could be that these anterior regions are especially subject to attentional modulations. This appears to be unique to when color is goal-relevant, as hV4 qualitatively responds just as strongly, if not stronger, to task-irrelevant but salient color stimuli than VO1/VO2 in a previous study (Thayer & Sprague, 2023). This intriguing functional distinction between extracting salient locations in posterior regions and relevant locations in anterior regions, which echoes similar sharp emergence of top-down cognitive signals in extrastriate cortex during visual working memory (Mendoza-Halliday et al., 2014), should be explored in future work spanning additional manipulations of visual stimuli and task demand conditions.

In sum, we found that feature-selective retinotopic regions in visual cortex track the location of stimuli throughout the visual field and showed a stronger response to those stimuli when each region's preferred feature was goal-relevant. Regions had different spatial profiles, with color regions locally enhancing activation at the stimulus location, and motion regions showing a global enhancement of activation. These findings emphasize the role of TO1/TO2 and hV4/VO1/VO2 as feature dimension maps and specify how they prioritize information in the visual field, which can be exploited to fully understand how we successfully navigate our visual environment.

References

Albright, T. D., & Stoner, G. R. (1995). Visual motion perception. *Proceedings of the National Academy of Sciences*, *92*(7), 2433–2440. https://doi.org/10.1073/pnas.92.7.2433

Amano, K., Wandell, B. A., & Dumoulin, S. O. (2009). Visual field maps, population receptive field sizes, and visual field coverage in the human MT+ complex. *Journal of Neurophysiology*, *102*(5), 2704–2718.

Awh, E., Belopolsky, A. V., & Theeuwes, J. (2012). Top-down versus bottom-up attentional control: A failed theoretical dichotomy. *Trends in Cognitive Sciences*, *16*(8), 437–443.

Bahle, B., Thayer, D. D., Mordkoff, J. T., & Hollingworth, A. (2019). The architecture of working memory: Features from multiple remembered objects produce parallel, coactive guidance of attention in visual search. *Journal of Experimental Psychology: General*.

Baker, D. H., Vilidaite, G., Lygo, F. A., Smith, A. K., Flack, T. R., Gouws, A. D., & Andrews, T. J. (2021). Power contours: Optimising sample size and precision in experimental psychology and human neuroscience. *Psychological Methods*, *26*, 295–314. https://doi.org/10.1037/met0000337

Beauchamp, M. S., Cox, R. W., & Deyoe, E. A. (1997). Graded Effects of Spatial and Featural Attention on Human Area MT and Associated Motion Processing Areas. *Journal of Neurophysiology*, *78*(1), 516–520. https://doi.org/10.1152/jn.1997.78.1.516

Benjamini, Y., & Yekutieli, D. (2001). The control of the false discovery rate in multiple testing under dependency. *The Annals of Statistics*, *29*(4), 1165–1188. https://doi.org/10.1214/aos/1013699998

Borst, A., & Egelhaaf, M. (1989). Principles of visual motion detection. *Trends in Neurosciences*, *12*(8), 297–306. https://doi.org/10.1016/0166-2236(89)90010-6

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*(4), 433–436. https://doi.org/10.1163/156856897X00357

Brewer, A. A., Liu, J., Wade, A. R., & Wandell, B. A. (2005). Visual field maps and stimulus selectivity in human ventral occipital cortex. *Nature Neuroscience*, *8*(8), 1102–1109.

Carrasco, M. (2011). Visual attention: The past 25 years. *Vision Research*, *51*(13), 1484–1525. https://doi.org/10.1016/j.visres.2011.04.012

Chawla, D., Rees, G., & Friston, K. J. (1999). The physiological basis of attentional modulation in extrastriate visual areas. *Nature Neuroscience*, *2*(7), 671–676. https://doi.org/10.1038/10230

Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, *3*(3), 201–215. https://doi.org/10.1038/nrn755

Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, *18*(1), 193–222.

Dumoulin, S. O., & Wandell, B. A. (2008). Population receptive field estimates in human visual cortex. *Neuroimage*, *39*(2), 647–660.

Eckstein, M. P. (2011). Visual search: A retrospective. *Journal of Vision*, *11*(5), 14–14. https://doi.org/10.1167/11.5.14

Ester, E. F., Sutterer, D. W., Serences, J. T., & Awh, E. (2016). Feature-Selective Attentional Modulations in Human Frontoparietal Cortex. *Journal of Neuroscience*, *36*(31), 8188–8199. https://doi.org/10.1523/JNEUROSCI.3935-15.2016

Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, *39*(2), 175–191. https://doi.org/10.3758/BF03193146

Folk, C. L., Remington, R. W., & Johnston, J. C. (1992). Involuntary covert orienting is contingent on attentional control settings. *Journal of Experimental Psychology: Human Perception and Performance*, *18*, 1030–1044. https://doi.org/10.1037/0096-1523.18.4.1030

Found, A., & Müller, H. J. (1996). Searching for unknown feature targets on more than one dimension: Investigating a "dimension-weighting" account. *Perception & Psychophysics*, *58*(1), 88–101. https://doi.org/10.3758/BF03205479

Harrison, A., Thayer, D., & Sprague, T. (2024). Stimulus representations in neural priority maps are equally enhanced by attention independent of the number of attended locations. *Journal of Vision*, *24*(10), 1507. https://doi.org/10.1167/jov.24.10.1507

Huk, A. C., Dougherty, R. F., & Heeger, D. J. (2002). Retinotopy and functional subdivision of human areas MT and MST. *Journal of Neuroscience*, *22*(16), 7195–7205.

Itti, L., & Koch, C. (2001). Computational modelling of visual attention. *Nature Reviews Neuroscience*, *2*(3), 194.

Kay, K. N., Winawer, J., Mezer, A., & Wandell, B. A. (2013). Compressive spatial summation in human visual cortex. *Journal of Neurophysiology*, *110*(2), 481–494.

Kriegeskorte, N., & Bandettini, P. (2007). Analyzing for information, not activation, to exploit high-resolution fMRI. *NeuroImage*, *38*(4), 649–662. https://doi.org/10.1016/j.neuroimage.2007.02.022

Leber, A. B., & Egeth, H. E. (2006). It's under control: Top-down search strategies can override attentional capture. *Psychonomic Bulletin & Review*, *13*(1), 132–138. https://doi.org/10.3758/BF03193824

Liesefeld, H. R., Liesefeld, A. M., & Müller, H. J. (2019). Distractor-interference reduction is dimensionally constrained. *Visual Cognition*, *27*, 247–259. https://doi.org/10.1080/13506285.2018.1561568

Liesefeld, H. R., Liesefeld, A. M., Pollmann, S., & Müller, H. J. (2019). Biasing Allocations of Attention via Selective Weighting of Saliency Signals: Behavioral and Neuroimaging Evidence for the Dimension-Weighting Account. In T. Hodgson (Ed.), *Processes of Visuospatial Attention and Working Memory* (pp. 87–113). Springer International Publishing. https://doi.org/10.1007/7854_2018_75

Liesefeld, H. R., & Müller, H. J. (2019). Distractor handling via dimension weighting. *Current Opinion in Psychology*, *29*, 160–167. https://doi.org/10.1016/j.copsyc.2019.03.003

Liesefeld, H. R., & Müller, H. J. (2023). *Target Salience and Search Modes: A Commentary on Theeuwes (2023)* (1). *6*(1), Article 1. https://doi.org/10.5334/joc.279

Luck, S. J., Gaspelin, N., Folk, C. L., Remington, R. W., & Theeuwes, J. (2020). Progress toward resolving the attentional capture debate. *Visual Cognition*. https://www.tandfonline.com/doi/full/10.1080/13506285.2020.1848949

Mackey, W. E., Winawer, J., & Curtis, C. E. (2017). Visual field map clusters in human frontoparietal cortex. *Elife*, *6*, e22974.

Mante, V., Sussillo, D., Shenoy, K. V., & Newsome, W. T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature*, *503*(7474), Article 7474. https://doi.org/10.1038/nature12742

Martínez-Trujillo, J. C., & Treue, S. (2002). Attentional Modulation Strength in Cortical Area MT Depends on Stimulus Contrast. *Neuron*, *35*(2), 365–370. https://doi.org/10.1016/S0896-6273(02)00778-X

Maunsell, J. H., & Treue, S. (2006). Feature-based attention in visual cortex. *Trends in Neurosciences*, *29*(6), 317–322.

McAdams, C. J., & Maunsell, J. H. R. (2000). Attention to Both Space and Feature Modulates Neuronal Responses in Macaque Area V4. *Journal of Neurophysiology*, *83*(3), 1751–1755. https://doi.org/10.1152/jn.2000.83.3.1751

McMains, S. A., Fehd, H. M., Emmanouil, T.-A., & Kastner, S. (2007). Mechanisms of feature-and space-based attention: Response modulation and baseline increases. *Journal of Neurophysiology*, *98*(4), 2110–2121.

Mendoza-Halliday, D., Torres, S., & Martinez-Trujillo, J. C. (2014). Sharp emergence of feature-selective sustained activity along the dorsal visual pathway. *Nature Neuroscience*, *17*(9), 1255–1262. https://doi.org/10.1038/nn.3785

Moore, T., & Armstrong, K. M. (2003). Selective gating of visual signals by microstimulation of frontal cortex. *Nature*, *421*(6921), 370–373.

Müller, H. J., Heller, D., & Ziegler, J. (1995). Visual search for singleton feature targets within and across feature dimensions. *Perception & Psychophysics*, *57*(1), 1–17. https://doi.org/10.3758/BF03211845

O'Craven, K. M., Downing, P. E., & Kanwisher, N. (1999). fMRI evidence for objects as the units of attentional selection. *Nature*, *401*(6753), 584–587.

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*(4), 437–442. https://doi.org/10.1163/156856897X00366

Poltoratski, S., Ling, S., McCormack, D., & Tong, F. (2017). Characterizing the effects of feature salience and top-down attention in the early visual system. *Journal of Neurophysiology*, *118*(1), 564–573. https://doi.org/10.1152/jn.00924.2016

Runeson, E., Boynton, G. M., & Murray, S. O. (2013). Effects of task and attentional selection on responses in human visual cortex. *Journal of Neurophysiology*, *109*(10), 2606–2617.

Saenz, M., Buracas, G. T., & Boynton, G. M. (2002). Global effects of feature-based attention in human visual cortex. *Nature Neuroscience*, *5*(7), 631–632.

Sauter, M., Liesefeld, H. R., Zehetleitner, M., & Müller, H. J. (2018). Region-based shielding of visual search from salient distractors: Target detection is impaired with same- but not different-dimension distractors. *Attention, Perception, & Psychophysics*, *80*(3), 622–642. https://doi.org/10.3758/s13414-017-1477-4

Schneider, K. A. (2011). Subcortical Mechanisms of Feature-Based Attention. *Journal of Neuroscience*, *31*(23), 8643–8653. https://doi.org/10.1523/JNEUROSCI.6274-10.2011

Schoenfeld, M., Hopf, J.-M., Martinez, A., Mai, H., Sattler, C., Gasde, A., Heinze, H.-J., & Hillyard, S. (2007). Spatio-temporal Analysis of Feature-Based Attention. *Cerebral Cortex*, *17*(10), 2468–2477. https://doi.org/10.1093/cercor/bhl154

Scolari, M., Ester, E. F., & Serences, J. T. (2014). Feature- and Object-Based Attentional Modulation in the Human Visual System. In A. C. (Kia) Nobre & S. Kastner (Eds.), *The Oxford Handbook*

of Attention (p. 0). Oxford University Press. https://doi.org/10.1093/oxfordhb/9780199675111.013.009

Serences, J. T., & Boynton, G. M. (2007). Feature-Based Attentional Modulations in the Absence of Direct Visual Stimulation. *Neuron*, *55*(2), 301–312. https://doi.org/10.1016/j.neuron.2007.06.015

Serences, J. T., & Saproo, S. (2012). Computational advances towards linking BOLD and behavior. *Neuropsychologia*, *50*(4), 435–446.

Serences, J. T., & Yantis, S. (2006). Selective visual attention and perceptual coherence. *Trends in Cognitive Sciences*, *10*(1), 38–45.

Sprague, T. C., Adam, K. C., Foster, J. J., Rahmati, M., Sutterer, D. W., & Vo, V. A. (2018). Inverted encoding models assay population-level stimulus representations, not single-unit neural tuning. *eNeuro*, *5*(3).

Sprague, T. C., Boynton, G. M., & Serences, J. T. (2019). *Inverted encoding models estimate sensible channel responses for sensible models*. https://doi.org/10.1101/642710

Sprague, T. C., Ester, E. F., & Serences, J. T. (2016). Restoring latent visual working memory representations in human cortex. *Neuron*, *91*(3), 694–707.

Sprague, T. C., Itthipuripat, S., Vo, V. A., & Serences, J. T. (2018). Dissociable signatures of visual salience and behavioral relevance across attentional priority maps in human cortex. *Journal of Neurophysiology*, *119*(6), 2153–2165.

Sprague, T. C., & Serences, J. T. (2013). Attention modulates spatial priority maps in the human occipital, parietal and frontal cortices. *Nature Neuroscience*, *16*(12), 1879.

Swisher, J. D., Halko, M. A., Merabet, L. B., McMains, S. A., & Somers, D. C. (2007). Visual Topography of Human Intraparietal Sulcus. *The Journal of Neuroscience*, *27*(20), 5326–5337. https://doi.org/10.1523/JNEUROSCI.0991-07.2007

Thayer, D. D., & Sprague, T. C. (2023). *Feature-specific salience maps in human cortex* (p. 2023.03.29.534828). bioRxiv. https://doi.org/10.1101/2023.03.29.534828

Theeuwes, J. (2010). Top–down and bottom–up control of visual selection. *Acta Psychologica*, *135*(2), 77–99. https://doi.org/10.1016/j.actpsy.2010.02.006

Theeuwes, J. (2023). *The Attentional Capture Debate: When Can We Avoid Salient Distractors and When Not?* (1). *6*(1), Article 1. https://doi.org/10.5334/joc.251

Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, *12*(1), 97–136.

Treue, S., & Trujillo, J. C. M. (1999). Feature-based attention influences motion processing gain in macaque visual cortex. *Nature*, *399*(6736), 575–579.

Wandell, B. A., Dumoulin, S. O., & Brewer, A. A. (2007). Visual field maps in human cortex. *Neuron*, *56*(2), 366–383.

Winawer, J., & Witthoft, N. (2015). Human V4 and ventral occipital retinotopic maps. *Visual Neuroscience*, *32*, E020. https://doi.org/10.1017/S0952523815000176

Witkowski, P. P., & Geng, J. J. (2023). Prefrontal Cortex Codes Representations of Target Identity and Feature Uncertainty. *Journal of Neuroscience*, *43*(50), 8769–8776. https://doi.org/10.1523/JNEUROSCI.1117-23.2023

Wolfe, J. M. (1994). Guided search 2.0 a revised model of visual search. *Psychonomic Bulletin & Review*, *1*(2), 202–238.

Wolfe, J. M., & Horowitz, T. S. (2004). What attributes guide the deployment of visual attention and how do they do it? *Nature Reviews Neuroscience*, *5*(6), 495.

Wolfe, J. M., & Horowitz, T. S. (2017). Five factors that guide attention in visual search. *Nature Human Behaviour*, *1*(3), 1–8. https://doi.org/10.1038/s41562-017-0058

Zhang, W., & Luck, S. J. (2009). Feature-based attention modulates feedforward visual processing. *Nature Neuroscience*, *12*(1), 24–25. https://doi.org/10.1038/nn.2223