Spatial selectivity in the temporoparietal junction, inferior frontal sulcus, and inferior parietal lobule

Kathleen A. Hansen
Laboratory of Brain and Cognition, National Institute of Mental Health, National Institutes of Health, Bethesda, MD, USA

Carlton Chu
Laboratory of Brain and Cognition, National Institute of Mental Health, National Institutes of Health, Bethesda, MD, USA

Annelise Dickinson
Laboratory of Brain and Cognition, National Institute of Mental Health, National Institutes of Health, Bethesda, MD, USA
Present address: George Mason University, Fairfax, VA, USA

Brandon Pye
Laboratory of Brain and Cognition, National Institute of Mental Health, National Institutes of Health, Bethesda, MD, USA
Present address: Georgia Tech University, Atlanta, GA, USA

J. Patrick Weller
Laboratory of Brain and Cognition, National Institute of Mental Health, National Institutes of Health, Bethesda, MD, USA
Present address: University of Washington, Seattle, WA, USA

Leslie G. Ungerleider
Laboratory of Brain and Cognition, National Institute of Mental Health, National Institutes of Health, Bethesda, MD, USA

Spatial selectivity, as measured by functional magnetic resonance imaging (fMRI) activity patterns that vary consistently with the location of visual stimuli, has been documented in many human brain regions, notably the occipital visual cortex and the frontal and parietal regions that are active during endogenous, goal-directed attention. We hypothesized that spatial selectivity also exists in regions that are active during exogenous, stimulus-driven attention. To test this hypothesis, we acquired fMRI data while subjects maintained passive fixation. At jittered time intervals, a briefly presented wedge-shaped array of rapidly expanding circles appeared at one of three contralateral or one of three ipsilateral locations. Positive fMRI activations were identified in multiple brain regions commonly associated with exogenous attention, including the temporoparietal junction, the inferior parietal lobule, and the inferior frontal sulcus. These activations were not organized as a map across the cortical surface. However, multivoxel pattern analysis of the fMRI activity correctly classified every pair of stimulus locations, demonstrating that patterns of fMRI activity were correlated with spatial location. These observations held for both contralateral and ipsilateral stimulus pairs as well as for stimuli of different textures (radial checkerboard) and shapes (squares and rings). Permutation testing verified that the obtained accuracies were not due to systematic biases and demonstrated that the findings were statistically significant.

Citation: Hansen, K. A., Chu, C., Dickinson, A., Pye, B., Weller, J. P., & Ungerleider, L. G. (2015). Spatial selectivity in the temporoparietal junction, inferior frontal sulcus, and inferior parietal lobule. *Journal of Vision, 15*(13):15, 1–15, doi:10.1167/15.13.15

doi: 10.1167/15.13.15 Received December 22, 2011; published September 18, 2015 ISSN 1534-7362

Introduction

Endogenous attention is attention that is goal directed and voluntary. Many cortical regions implicated in endogenous attention—including the intraparietal sulcus (IPS), frontal eye field, and inferior frontal sulcus (IFS)—are spatially selective. That is, the overall pattern of activity depends on a spatial factor such as the location of covert attention (Hansen, Kay, & Gallant, 2007; Silver & Kastner, 2009; Silver, Ress, & Heeger, 2005) or the destination of a saccadic eye movement (Hagler, Riecke, & Sereno, 2007; Kastner & McCains, 2007; Konen & Kastner, 2008; Levy, Schluppeck, Heeger, & Glimcher, 2007; Schluppeck, Glimcher, & Heeger, 2005; Sereno, Pitzalis, & Martinez, 2001; Silver & Kastner, 2009). In some cases, spatial selectivity can be detected experimentally even during passive fixation, although signals are typically more reliable when subjects are performing spatial attention tasks (Bressler & Silver, 2010).

These experimental demonstrations of spatial selectivity imply that the underlying neural populations have access to information about spatial location. Hypothetically, spatial information might allow a neural population to focus available computing resources on stimuli in the relevant parts of visual space. It is reasonable to assume that access to spatial information would also be adaptive for regions implicated in exogenous attention—attention that is stimulus driven and involuntary. Therefore, it is striking that spatial selectivity has not previously been demonstrated in certain regions commonly associated with exogenous attention, namely the temporoparietal junction (TPJ) and inferior parietal lobule (IPL; Arrington, Carr, Mayer, & Rao, 2000; Corbetta & Shulman, 2002; Dodds, Morein-Zamir, & Robbins, 2011; Downar, Crawley, Mikulis, & Davis, 2000; Macaluso, Frith, & Driver, 2002; Shulman et al., 2010).

We hypothesized that previous efforts to reveal spatial selectivity in these regions failed not because there is no spatial selectivity but rather because the cortical locations selective for different parts of space are overlapping and distributed (Cichy, Chen, & Haynes, 2011; Hassabis et al., 2009). To detect the hypothetical spatial selectivity, therefore, we used multivoxel pattern analysis (MVPA) on functional magnetic resonance imaging (fMRI) data acquired while fixating subjects passively viewed brief, infrequently presented stimuli at various spatial locations. Unlike traditional analysis methods, MVPA can detect fMRI selectivity even when the functional anatomy is overlapping and distributed (Carlson, Schrater, & He, 2003; Cichy et al., 2011; Cox & Savoy, 2003; Hanson, Matsuka, & Haxby, 2004; Haxby et al., 2001; Haynes & Rees, 2006; Kamitani & Tong, 2005; Kriegeskorte & Bandettini, 2007; Liu, Hospadaruk, Zhu, & Gardner, 2011; Norman, Polyn, Detre, & Haxby, 2006).

The MVPA approach measures how accurately a classifier algorithm can use fMRI data points from a given region of interest (ROI) to predict which experimental condition was present when the data points were acquired (Mur, Bandettini, & Kriegeskorte, 2009). Accurate predictions imply that patterns of activity across the ROI depend on experimental condition. Our MVPA results indicated that activity in TPJ, IFS, and IPL depended on the spatial locations of visual stimuli. Permutation testing (Nichols & Holmes, 2001) and subsequent group-level significance testing verified that these findings were statistically significant and could not be ascribed to systematic biases in the data.

Method

Twelve subjects (seven females, five males; ages 23–54 years) were studied. They were right handed and had normal color vision and normal or corrected-to-normal visual acuity. All subjects gave informed consent according to the guidelines set forth by the Laboratory of Brain and Cognition. All procedures were approved by the National Institute of Mental Health Institutional Review Board.

In Experiment 1, stimuli (Figure 1) were one of six adjacent wedge-shaped texture patterns presented on a grayscale background. The texture was an array of black dots whose radii increased rapidly to create a "looming" appearance of approaching the observer. This dynamic texture had several virtues. First, briefly presented, expanding circles are known to elicit exogenous attention (Tse, 2010), similar to other types of sudden motion (Jonides & Yantis, 1988). Second, the texture's presentation did not produce an afterimage. Third, unlike the more traditional radial checkerboard texture used in Experiment 2 (which reversed contrast every 125 ms for a total presentation time of 500 ms), the short frame duration permitted each expanding dot stimulus presentation to be completed in only 132 ms. This is not enough time to make a saccade. Thus, in Experiment 1, even if a subject broke fixation when a wedge appeared, the wedge would not have been present when the eye movement was completed. Each wedge subtended a 60° angle, was entirely within either the left or the right visual field, and extended to approximately ±5 radial degrees from the central fixation point. The wedges were presented one at a time. Each of the six wedges was presented six times per run, and five runs were collected. Each wedge presentation consisted of four frames shown for 33 ms each, with an average interval of 11.25 s between

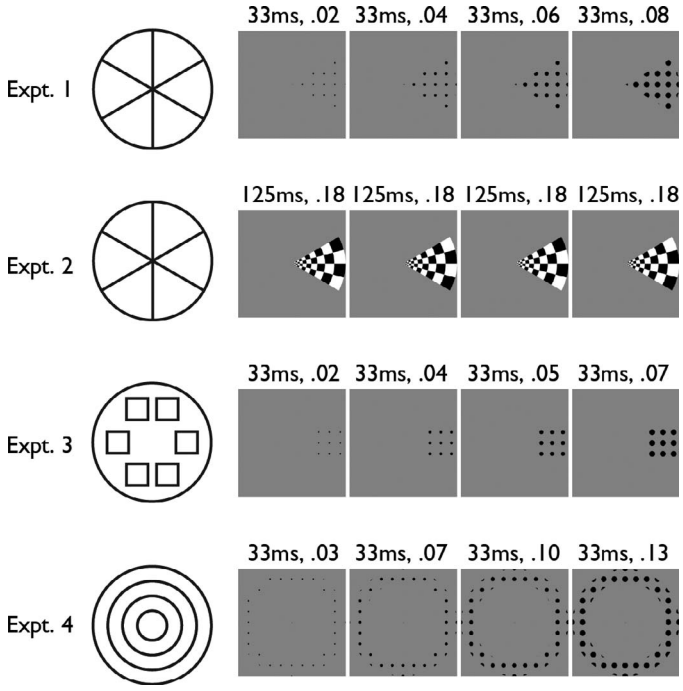

Figure 1. Stimuli. Left: Schematics showing the stimulus locations. The lines pictured in the schematics did not appear in the stimuli. Right: Example stimulus presentations. Numbers indicate each frame's duration and root mean square contrast.

presentation onsets ($SD = 2.5$ s; minimum = 7.5 s; maximum = 15 s). All 12 subjects participated in Experiment 1. In all experiments, the subjects maintained central fixation and passively viewed the eccentrically presented stimuli.

In Experiment 2, stimuli (Figure 1) were one of six adjacent wedge-shaped texture patterns presented on a grayscale background. The spatial locations and extents of the wedges were the same as those used in Experiment 1. The texture was black-and-white reversing-contrast radial checkerboard. The wedges were presented one at a time. Each of the six wedges was presented six times per run, and four or five runs were collected. Each wedge presentation consisted of four frames shown for 125 ms each (8-Hz reversal), with an average interval of 11.25 s between presentation onsets ($SD = 2.5$ s; minimum = 7.5 s; maximum = 15 s). Ten of the subjects participated in Experiment 2.

In Experiment 3, stimuli (Figure 1) were one of six square-shaped texture patterns presented on a grayscale background. The texture was the same expanding dot array used in Experiment 1. The square texture patterns were identical to one another, varying only in position via translation. The center of each square was inside the space occupied by one of the Experiment 1 wedges, and no part of any square extended outside that wedge's space. The squares were presented one at a time. Each of the six squares was presented six times per run, and five runs were collected. Each square presentation consisted of four frames shown for 33 ms each, with an average interval of 11.25 s between presentation onsets ($SD = 2.5$ s; minimum = 7.5 s; maximum = 15 s). Four of the subjects participated in Experiment 3.

In Experiment 4, stimuli (Figure 1) were one of four adjacent, concentric, ring-shaped texture patterns presented on a grayscale background. The rings were of uniform width. The set of all four rings covered the same circular portion of the visual field that was covered by the set of all six wedges used in Experiments 1 and 2. The texture was the same expanding dot array used in Experiment 1. The rings were presented one at a time. Each of the four rings was presented eight times per run, and four runs were collected. Each ring presentation consisted of four frames shown for 33 ms each, with an average interval of 11.5 s between presentation onsets ($SD = 2.8$ s; minimum = 7.5 s; maximum = 15 s). Nine of the subjects participated in Experiment 4.

All magnetic resonance imaging (MRI) data were collected on a GE (Fairfield, CT) 3-Tesla scanner with a GE whole-head eight-channel coil. For fMRI we used an echo-planar imaging sequence with a repetition time of 2.5 s per shot (2.5 s per acquired brain volume), echo time of 30 ms, field of view of 22 cm × 22 cm, resolution of 64 × 64 voxels per slice (in-plane voxel size = 3.4 mm × 3.4 mm), and slice thickness of 3.0 mm. Each fMRI brain volume consisted of 38 axial slices. For anatomical images we used a magnetization prepared rapid acquisition gradient echo sequence with a field of view of 24 cm × 24 cm, 128 locations per slab, and slice thickness of 1.2 mm. Preprocessing and subsequent analysis of the MRI and fMRI data were performed with the AFNI and Freesurfer software packages (Cox, 1996; Cox & Hyde, 1997; Dale et al., 1999; Fischl, Sereno, & Dale, 1999; Fischl, Sereno, Tottell, & Dale, 1999). The anatomical images were used to create virtual cortical surfaces for each subject. For each experiment, the cortical surfaces of the subjects participating in that experiment were averaged to create a group cortical surface. The fMRI volumes for each subject were registered to the anatomical volume used to create that subject's surface representation via the *bbregister* function in Freesurfer. The fMRI data were projected onto the surface representation via the *mri_vol2surf* function in Freesurfer.

The ROIs for MVPA on the Experiment 1, 2, and 3 data were derived from the Experiment 4 data. The ROIs for MVPA on the Experiment 4 data were derived from the Experiment 1 data. The procedure for ROI identification was the same in both cases. We used a general linear model (GLM) to estimate each subject's fMRI responses to stimuli presented at each spatial location and counted the number of spatial locations evoking an fMRI response of $p < 0.01$

(uncorrected) at each voxel. We averaged these numbers across subjects on the average cortical surface and isolated clusters of contiguous nodes with positive values. This process ensured that the ROIs represented brain areas where stimuli evoked positive fMRI activations, which presumably reflect overall engagement of mechanisms responsive to the stimuli (Jimura & Poldrack, 2012). If we had not ensured that every ROI responded positively to at least one stimulus, it might have been difficult to interpret the MVPA results, which could conceivably have been driven exclusively by negative activations. Such a scenario could be problematic because negative activations often arise in brain regions where the function is not relevant to the experimental condition being tested. For example, when task-relevant stimuli are visual, stimulus presentation often results in negative activation of the auditory cortex (Amedi et al., 2005, Haxby et al., 1994). To create ROIs for each individual subject, the central coordinates of each surviving cluster were first identified. Volumetric spherical masks with an 18-mm radius (representing approximately 700 voxels) were created at these central coordinates and converted to each individual's surface space. The use of fixed-radius spherical masks ensured that the ROIs used for MVPA were of a consistent size.

To obtain input features for the MVPA, we used a GLM to estimate each subject's fMRI responses to every individual stimulus presentation. Each GLM represented one experimental run, and each regressor in a GLM represented the convolution of the timing of a single stimulus presentation with a gamma function. For example, the GLM for run 1 of Experiment 1 included 36 regressors (one for each individual wedge presentation). The GLMs also included motion estimates as regressors of no interest. In this article, the term *sample* refers to one such estimate of a subject's fMRI response to an individual stimulus presentation.

To perform the MVPA, within each ROI the N samples for each condition pair in an experiment—for example, the N samples for presentations of squares on the upper and lower right—were conceptually represented as points in a D-dimensional space, where D is the number of features (nodes in the ROI). Support vector machines identified the hyperplane that optimally divided the points into groups by maximizing the margin between two groups (Vapnik, 1995); in this example, the groups would be upper right versus lower right. The term *training* refers to this process. During training, one sample was left out (the leave one out cross-validation method). The sign of the distance between the left-out sample and the hyperplane was used to determine the class of the left-out data. The prediction was compared with the real left-out data to yield a single binary measurement: Was the prediction correct or not? During cross-validation, the whole procedure was repeated N times such that every sample was left out once, thus obtaining many binary measurements of prediction accuracy. Within each subject the ratio of correct to total predictions was calculated, producing an individual percentage correct measurement. The individual percentage correct measurements were averaged across subjects to obtain an overall percentage correct for (in this example) upper right versus lower right selectivity.

Results

ROI identification

The Experiment 1 and 4 data sets both produced multiple clusters of activations to at least one spatial location (see Method for the procedure for identifying clusters and converting them to ROIs). In each case, one very large cluster covered much of the occipital lobe and extended into the IPS bilaterally and into the TPJ in the right hemisphere. Spatial selectivity in the activated occipital cortex and IPS has already been well documented in retinotopic mapping studies (Silver & Kastner, 2009), but spatial selectivity in the right TPJ has not. Therefore, we used an anatomical landmark—the fundus of the superior temporal sulcus—to separate the large activated region into a TPJ cluster and a non-TPJ cluster. Within the fundus, all cortex anterior to the line running along the deepest part of the sulcus was assigned to TPJ, and all cortex posterior to the line was excluded from TPJ. The non-TPJ cluster and all clusters of less than 1000 mm^2 were discarded. Surviving clusters obtained from the Experiment 4 data (Figure 2; Table 1) were located in the right superior frontal gyrus (SFG), bilateral IFS, right TPJ, left IPL, and left fusiform gyrus (FG). Surviving clusters obtained from the Experiment 1 data (Figure 2; Table 2) were located in the right IFS, right TPJ, and left IPL. Thus, both Experiment 1 and Experiment 4 produced clusters in the right TPJ, right IFS, and left IPL.

The process used to convert clusters to ROIs ensured that the ROIs of the left FG, left IPL, right TPJ, right IFS, left IFS, and right SFG were the same size. We were concerned, however, that using the same process in the occipital cortex and the parietal cortex could result in ROIs representing arbitrarily varying portions of retinotopic areas, which could complicate the interpretation of the spatial selectivity measurements. To simplify the presentation, therefore, we derived a single large visual cortex (VC) ROI (Figure 2) from anatomical landmarks. The landmarks were chosen to avoid any overlap with the other ROIs and

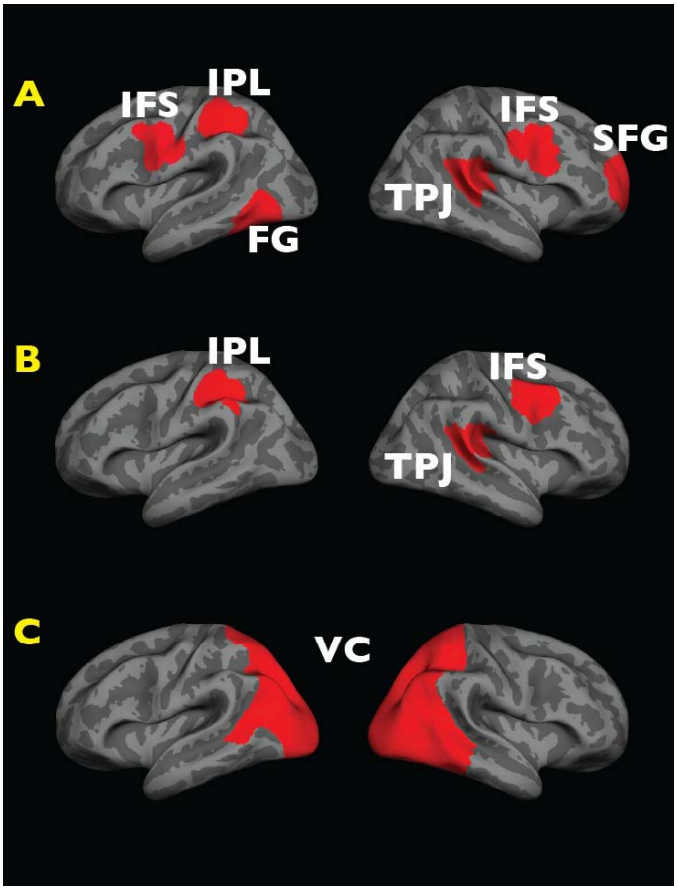

Figure 2. (A) ROIs derived from Experiment 4 data and used for MVPA analysis of data from Experiments 1 through 3. (B) ROIs derived from Experiment 1 data and used for MVPA analysis of data from Experiment 4. (C) VC ROI derived from anatomical landmarks and used for MVPA analysis of data from Experiments 1 through 4.

to avoid omitting any part of the occipital or parietal cortex where spatial selectivity was predicted based on previous observations of retinotopy. The retinotopic organization of many parts of the occipital and parietal cortex is still somewhat controversial; for simplicity, we chose to err on the side of inclusion. The resulting VC ROI might have included some regions that are not retinotopic or otherwise spatially selective, but this was not a great concern because nonpredicting data points would tend to be ignored during the classifier training. The VC ROI was converted to each individual's surface space for use with data from all four experiments.

Multivoxel pattern analysis

A pattern classifier was trained to predict which experimental condition was present when fMRI data points were acquired. Results are typically presented in percentage units that refer to the percentage of sample class predictions that are correct. (The term *samples* refers here to GLM estimates; for details, see Method.) The accuracy percentages are computed using cross-validation. High accuracy percentages imply that there are robust patterns of activity across the ROI that correlate with experimental condition.

Figures 3 through 6 present percentage correct numbers for each experimental condition pair in each experiment. For example, consider the data in Figure 3. For each condition pair (e.g., a wedge presented on the right at horizontal vs. a wedge presented on the lower right), the percentage correct classification is shown by the color of the corresponding matrix element. Chance performance would be 50% (green). Predictions for all condition pairs within all ROIs were substantially more accurate than chance in all experiments.

The matrix format in Figures 3 through 6 allows the viewer to identify systematic differences between condition pairs. For example, consider the results from the VC in the two wedge data sets (Experiments 1 and 2; Figures 3 and 4). The elements next to the matrix diagonal tend to be orange, while the elements away from the diagonal are red. The color difference implies that the VC data enabled the classifier to distinguish more easily between presentations of nonadjacent wedges (away from the diagonal) than adjacent wedges (next to the diagonal). This result is

ROI name	*x*	*y*	*z*
FG (LH)	−43.7	−58.6	−11.7
IPL (LH)	−37.0	−39.5	45.1
TPJ (RH)	46.5	−39.1	11.5
IFS (RH)	44.3	0.4	30.5
IFS (LH)	−53.9	0.3	34.2
SFG (RH)	22.6	57.8	11.1
VC	*	*	*

Table 1. Coordinates of ROIs used for MVPA analysis of Experiment 1 through 3 data. LH = left hemisphere; RH = right hemisphere. * The large VC ROI occupied most of the occipital lobe and part of the parietal lobe (see text). Coordinates in this and other tables are approximations of Talairach coordinates derived from a nonlinear transform of MNI305 coordinates by Freesurfer.

ROI name	*x*	*y*	*z*
IPL (LH)	−40.4	−39.3	39.3
TPJ (RH)	51.0	−38.6	10.2
IFS (RH)	45.3	3.7	42.7
VC	*	*	*

Table 2. Coordinates of ROIs used for MVPA analysis of Experiment 4 data. LH = xx; RH = xx. * The large VC ROI occupied most of the occipital lobe and part of the parietal lobe (see text).

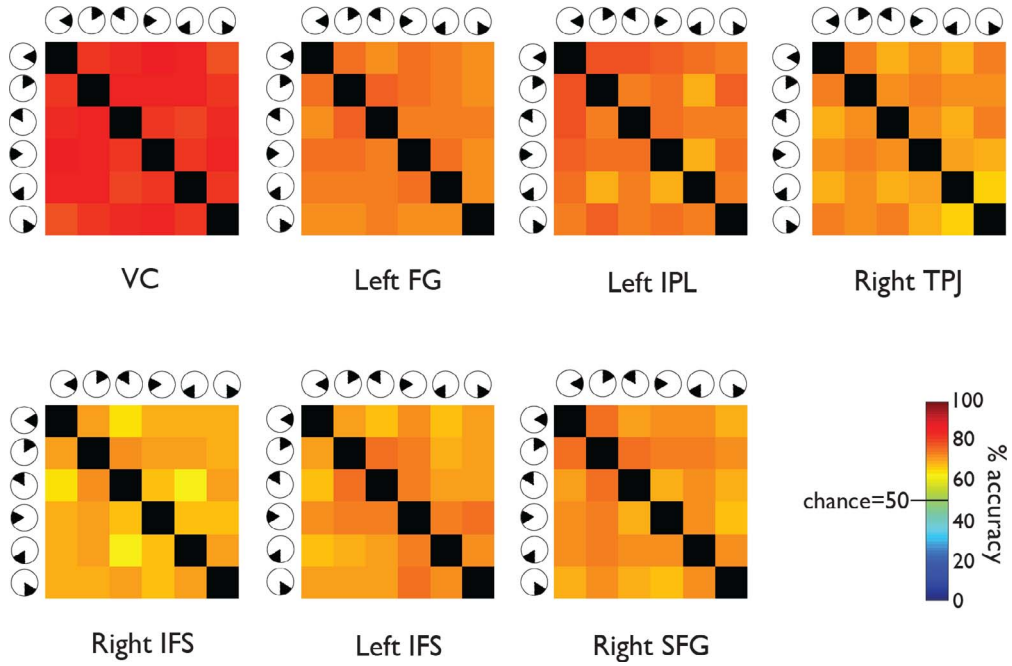

Figure 3. MVPA prediction accuracies, Experiment 1. Each matrix provides accuracy predictions derived from one ROI. The color of each element in a matrix represents the mean percentage correct, averaged across subjects, of a particular condition pair's MVPA predictions. For example, the element in the top right corner of each matrix represents the percentage of correct classifications of wedges presented on the right at horizontal versus on the lower right. Chance performance would be 50% (green). Predictions for all condition pairs were substantially more accurate than chance.

intuitive because representations of adjacent stimulus locations are retinotopically adjacent in the VC (e.g., Tootell, Silverman, Switkes, & De Valois, 1982). The absence of the same pattern in the squares data set (Experiment 3; Figure 5) is also intuitive because the squares, unlike the wedges, were not immediately adjacent to one another (Figure 1). In the ring data set (Experiment 4; Figure 6), the element in the upper right corner of the matrix is red, and the element next to the upper left corner of the matrix is light orange. The color difference implies that the VC data enabled the classifier to distinguish more easily between presentations of the innermost versus outermost rings (upper right corner) than between presentations of the innermost versus second innermost ring (next to the upper left corner). Like the wedges, the rings were adjacent to one another (Figure 1), so the observed pattern in the VC is expected. We looked for such patterns in the percentage correct matrices from ROIs outside the VC but did not observe them.

In all ROIs, the classifier was able to distinguish every contralateral wedge or square from every ipsilateral wedge or square, every contralateral wedge or square from every other contralateral wedge or square, and every ipsilateral wedge or square from every other ipsilateral wedge or square. Such findings are expected for the VC ROI, which included both hemispheres, but they are striking for the other ROIs, which represented data from one hemisphere only. The findings imply that patterns of activity in the left FG, left IPL, right TPJ, right IFS, left IFS, and right SFG regions correlated with stimulus presentations in both halves of space.

To investigate laterality for each ROI quantitatively, we used the MVPA results from Experiments 1 through 3 to perform a one-tailed t test of the hypothesis that accuracies were higher on average for predictions about condition pairs representing two contralateral wedges (or squares) compared with two ipsilateral wedges (or squares). The t test had 74 df; and input consisted of within-ROI accuracy percentages for every pair of contralateral or ipsilateral stimuli in Experiments 1, 2, and 3 for every ROI in every individual subject. (For this laterality analysis, we divided VC into separate left- and right-hemisphere ROIs.) A positive result would indicate that the classifier distinguished more easily between contralateral stimulus locations than between ipsilateral stimulus locations. Results were positive in the left VC ($p < 0.01$, $t = 3.183$), right VC ($p < 0.00001$, $t = 4.777$), and left IPL ($p < 0.001$, $t = 3.171$) and null in the remaining ROIs (Table 3; $p > 0.5$ in each ROI).

Note that despite the VC prediction advantage for contralateral versus contralateral over ipsilateral versus ipsilateral stimulus pairs, the ipsilateral versus ipsilateral predictions themselves were better than chance. In fact, the mean ipsilateral versus ipsilateral prediction in the VC exceeded 75%. This is a useful demonstration of

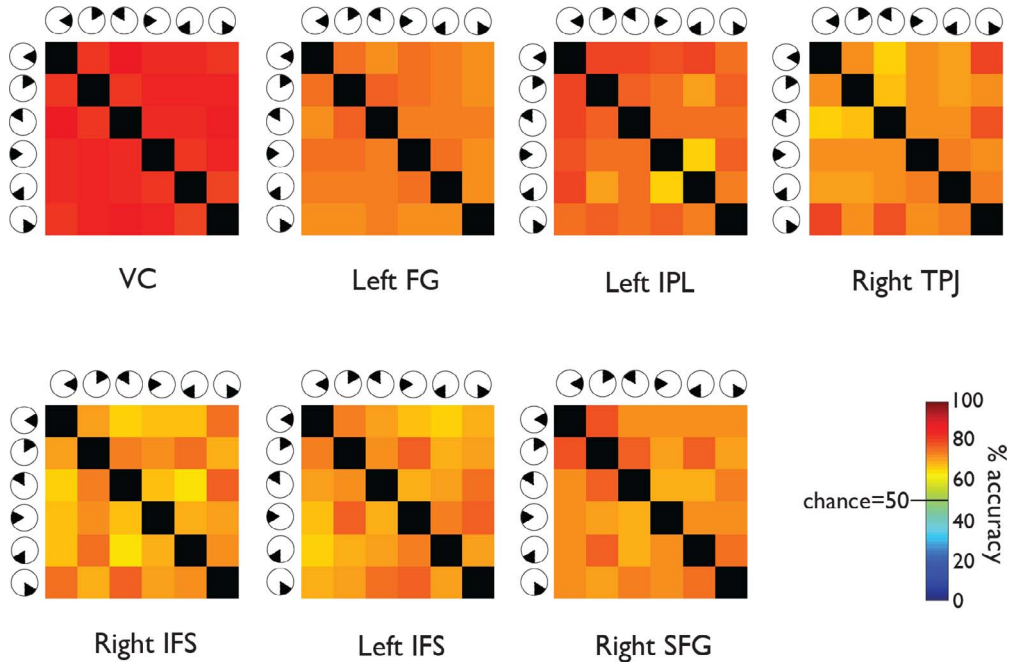

Figure 4. MVPA prediction accuracies, Experiment 2. Each matrix provides accuracy predictions derived from one ROI. The color of each element in a matrix represents the mean percentage correct, averaged across subjects, of a particular condition pair's MVPA predictions. For example, the element in the top right corner of each matrix represents the percentage of correct classifications of wedges presented on the right at horizontal versus on the lower right. Chance performance would be 50% (green). Predictions for all condition pairs were substantially more accurate than chance.

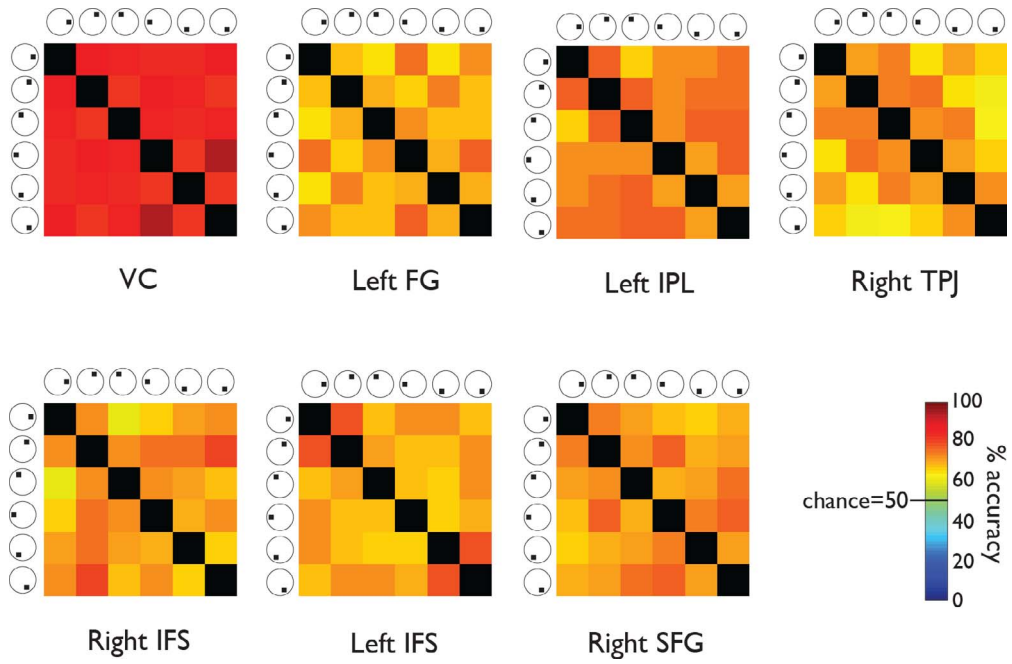

Figure 5. MVPA prediction accuracies, Experiment 3. Each matrix provides accuracy predictions derived from one ROI. The color of each element in a matrix represents the mean percentage correct, averaged across subjects, of a particular condition pair's MVPA predictions. For example, the element in the top right corner of each matrix represents the percentage of correct classifications of squares presented on the right at horizontal versus on the lower right. Chance performance would be 50% (green). Predictions for all condition pairs were substantially more accurate than chance.

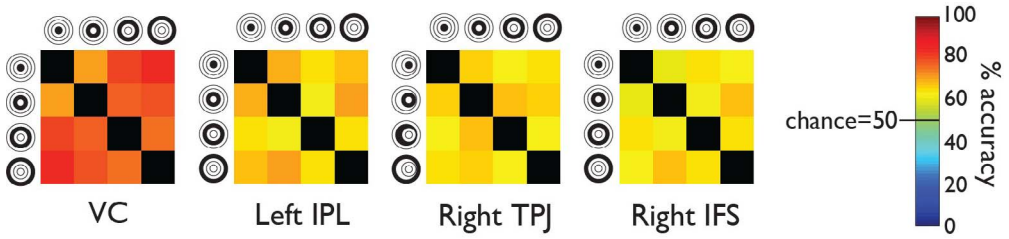

Figure 6. MVPA prediction accuracies, Experiment 4. Each matrix provides accuracy predictions derived from one ROI. The color of each element in a matrix represents the mean percentage correct, averaged across subjects, of a particular condition pair's MVPA predictions. For example, the element in the top right corner of each matrix represents the percentage of correct classifications of rings presented closest to versus farthest from the fixation point. Chance performance would be 50% (green). Predictions for all condition pairs were substantially more accurate than chance.

a point that is important to state clearly: The ability of an ROI to differentiate between ipsilateral (or contralateral) locations does not constitute evidence that every cortical location inside these ROIs is selective for different ipsilateral (or contralateral) locations. Rather, the ipsilateral (or contralateral) predictions constitute evidence that information about ipsilateral (or contralateral) locations is present in the set of data points across the ROI.

Figure 7 summarizes our main findings by presenting the accuracy percentages, averaged across subjects and condition pairs, for all ROIs and experiments. Chance performance would be 50%. In all cases the accuracy percentages were substantially greater than chance, demonstrating that patterns of activity in all ROIs depended on the spatial location of the stimulus. Because the four experiments used stimuli of different textures (expanding dots and radial checkerboard) and shapes (wedges, squares, and rings), the demonstration of spatial selectivity generalized across stimulus classes.

Permutation testing and group-level significance testing

We also performed nonparametric permutation testing (Nichols & Holmes, 2001) and subsequent group-level significance testing on the obtained accuracy percentages. In the permutation testing, for each training sample pair a 50% random choice was made of whether to switch the sample labels. The cross-validation was redone with the randomly labeled samples. This process was repeated 200 times to derive a distribution of accuracy percentages expected to occur by chance (i.e., the null distribution). The actual observed percentage was compared with the derived null distribution to yield a p value estimate for one condition pair prediction in one ROI in one subject during one experiment. We refer to this p value estimate as an *individual p value*. In every ROI and experiment, the most commonly obtained individual p value was 0.005. Note that a p value of 0.005 is not only low but is actually the lowest possible output of a cross-validation procedure with 200 repeats. Thus, very high levels of significance, even at the individual subject level, were the norm in our data set.

For group-level significance testing, the individual p values were converted to z scores. Simple one-sample t tests on the individual z scores, within experiment and ROI and across subjects and condition pairs, provided group p values. In every experiment and ROI, the group p values were less than 10^{-18} (Figure 7, asterisks), demonstrating that the findings were highly significant. Since the t tests were performed on the output of permutation testing with cross-validation, the

ROI	Experiment 1	Experiment 2	Experiments 1 through 3 combined
VC (RH)	<0.01, 2.579, 35	<0.001, 3.896, 26	<0.00001, 4.777, 74
VC (LH)	0.055, 1.640, 35	<0.01, 2.959, 26	<0.01, 3.183, 74
FG (LH)	ns (>0.2), −0.950, 35	ns (>0.2), −1.274, 26	ns (>0.5), −1.611, 74
IPL (LH)	<0.05, 2.152, 35	<0.01, 2.669, 26	<0.01, 3.171, 74
TPJ (RH)	ns (>0.2), −0.971, 35	ns (>0.2), −0.764, 26	ns (>0.5), −0.331, 74
IFS (RH)	ns (>0.2), −2.113, 35	ns (>0.2), −1.824, 26	ns (>0.5), −3.170, 74
IFS (LH)	ns (>0.2), −0.879, 35	ns (>0.2), −0.162, 26	ns (>0.5), −0.032, 74
SFG (RH)	ns (>0.2), −0.732, 35	ns (>0.2), −1.057, 26	ns (>0.5), −1.236, 74

Table 3. Contralaterality. Results (p, t, and df) of a one-tailed t test testing the hypothesis that the predictions of contralateral versus contralateral stimuli were more accurate than the predictions of ipsilateral versus ipsilateral stimuli.

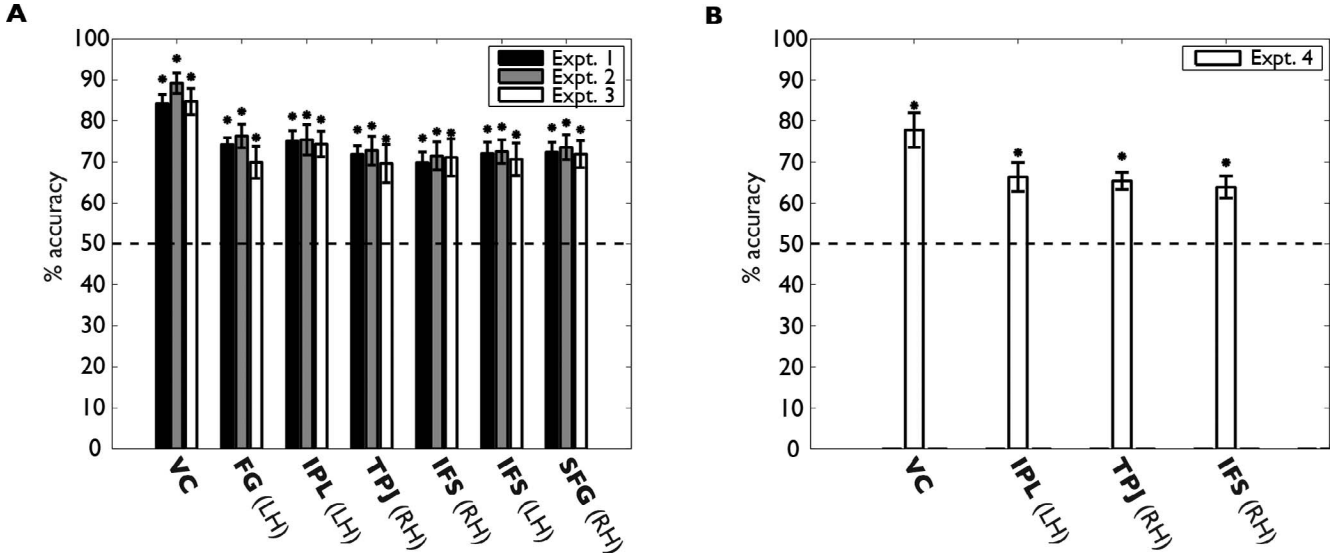

Figure 7. Mean MVPA prediction accuracies, all experiments. (A) Accuracies of predictions derived from the data in Experiments 1 through 3 within the ROIs derived from the Experiment 4 data. The height of each bar represents the mean percentage correct, averaged across subjects and condition pairs, of the MVPA predictions; error bars represent ± 1 *SD* across condition pairs. Chance performance would be 50%. In all cases prediction accuracy was significantly greater than chance. Asterisks indicate $p < 10^{-18}$; p values were obtained by performing a one-sample t test on the results of the within-subject, within-condition-pair permutation tests. (B) Accuracies of predictions derived from the data in Experiment 4 within the ROIs derived from the Experiment 1 data.

very low group p values also ruled out systematic bias in the data (e.g., bias caused by within-run trends in the fMRI signal) as an explanation for the finding that the observed prediction accuracies exceeded chance (Chen et al., 2011).

The prediction accuracies obtained in the anatomically defined VC ROI were higher than those obtained in the cluster-defined ROIs (Figures 3 through 6; compare red–orange in the VC with light orange–yellow elsewhere). This was not surprising because the VC ROI included the cortex already known to be organized retinotopically across the cortical surface. However, the higher prediction accuracies could conceivably have resulted from the fact that the VC ROI was considerably larger than the other ROIs (Figure 2) and so provided the MVPA with more training data. To check this possibility, we randomly sampled a subset of data points from the VC ROI such that the downsampled VC ROI and the other ROIs were the same size. Prediction accuracies in the downsampled VC ROI were still higher than those obtained in the cluster-defined ROIs (compare red–orange in Figure 8 with light orange–yellow in Figures 3 through 6).

Nonadjacent versus adjacent stimulus pairs

By inspection, the color matrices in Figures 3 through 6 show that, in the VC, the accuracy for decoding stimulus position from the pattern of blood-oxygen-level dependent (BOLD) activity was higher for nonadjacent stimuli than for adjacent stimuli. This result was expected from the retinotopic organization known to be present in the VC ROI. Conversely, this tendency was not observed in the other ROIs. The lack of a prediction advantage for nonadjacent over adjacent stimulus pairs outside VC implies either that selectivity for stimulus location in these regions is not organized retinotopically across the cortical surface or that any existing retinotopic organization is too coarse to be detectable in our data set. Given the importance of this observation, it seemed desirable to show it quantitatively. Figure 9 plots the difference in mean prediction accuracies between nonadjacent versus adjacent stimulus pairs across all four experiments; stars indicate a difference significantly ($p < 0.05$) greater than zero, as measured by a one-tailed t test across subjects and stimulus pairs. In the VC, but not in the remaining ROIs, nonadjacent stimulus pairs predicted better than adjacent stimulus pairs in Experiments 1, 2, and 4. (Note that in Experiment 3 none of the stimuli were literally adjacent, so a reduced prediction advantage in the VC for Experiment 3 relative to Experiment 1 is expected.) These results confirm the apparent absence of retinotopic organization for stimulus selectivity in the ROIs outside VC.

To further document the apparent lack of topography in the ROIs outside the VC, we converted each

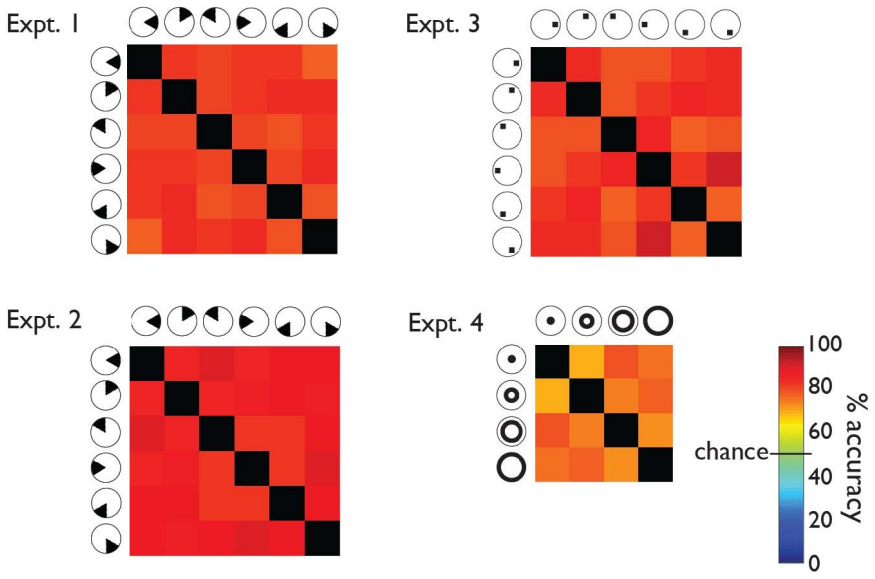

Figure 8. MVPA prediction accuracies in the downsampled VC ROI. A subset of data points was randomly sampled from the VC ROI such that the downsampled VC ROI and the other ROIs were the same size. In this figure, each matrix provides accuracy predictions derived from the downsampled VC ROI in one experiment. The color of each element in a matrix represents the mean percentage correct, averaged across subjects, of a particular condition pair's MVPA predictions. For example, the element in the top right corner of the Experiment 1 matrix represents the percentage of correct classifications of wedges presented on the right at horizontal versus on the lower right. Chance performance would be 50% (green). Predictions for all condition pairs were substantially more accurate than chance.

voxel's GLM activations to the various Experiment 1 (wedge) and Experiment 4 (ring) stimuli to estimates of preferred angle and eccentricity using the polar average for the angle and the center of mass for the eccentricity estimates. When these estimates were projected onto the cortical surface the familiar topographic maps in the early VC were apparent, but we could not detect topographic organization in the other ROIs. Surface

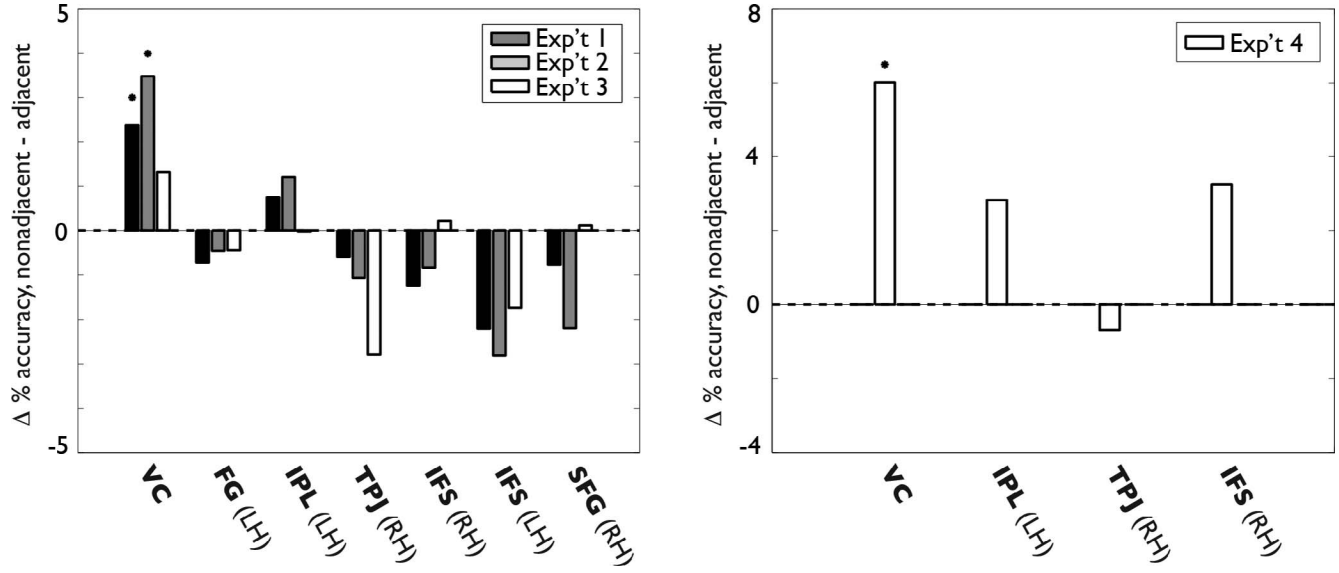

Figure 9. Mean MVPA prediction accuracy differences for nonadjacent versus adjacent stimulus pairs, all experiments. Stars indicate a difference significantly ($p < 0.05$) greater than zero as measured by a one-tailed t test across subjects and stimulus pairs. In the VC, but not in the remaining ROIs, nonadjacent stimulus pairs predicted better than adjacent stimulus pairs in Experiments 1, 2, and 4. (Note that in Experiment 3, none of the stimuli were literally adjacent, so a reduced prediction advantage in the VC for Experiment 3 relative to Experiment 1 is expected.)

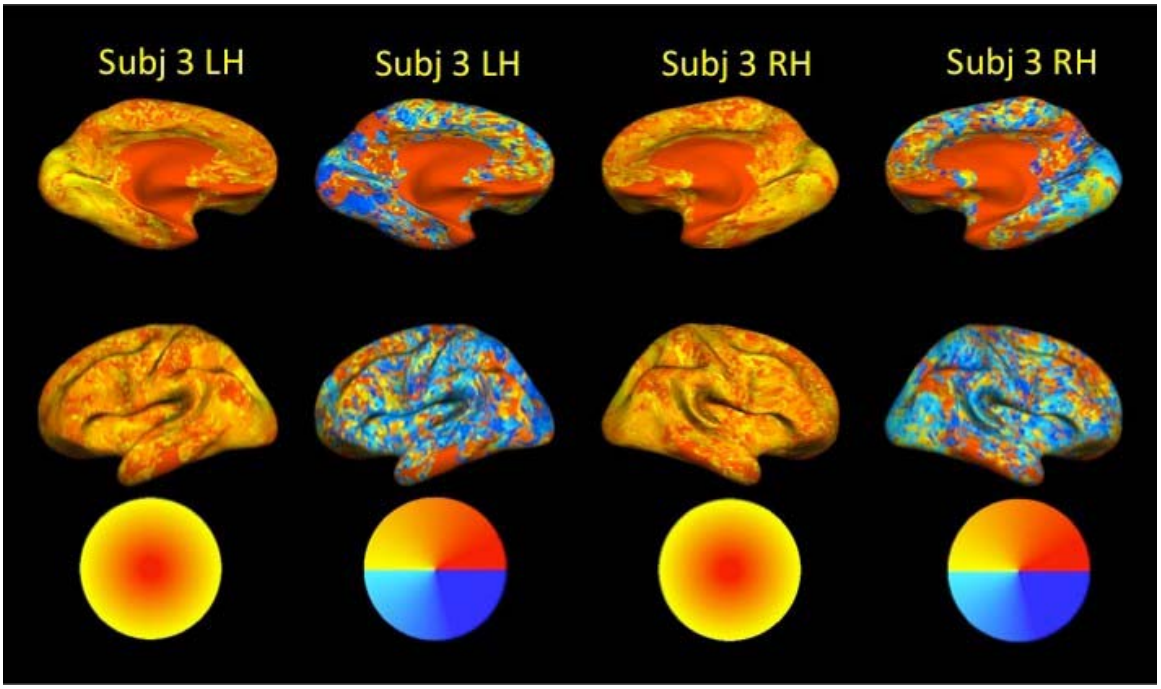

Figure 10. Voxelwise preferred angle and eccentricity estimates in an example subject. Preferred angles were calculated using the polar mean on the GLM activations from Experiment 1; preferred eccentricities were calculated using the center of mass on the GLM activations from Experiment 4. Despite the coarse spatial resolution used in our whole-brain design, topographic organization was evident in the early VC. We did not observe topographic organization in the other ROIs.

maps of preferred angle and eccentricity estimates from an example subject are shown in Figure 10. One caveat here is that in order to obtain whole-brain coverage in this study, we chose to use a sampling resolution coarser than what is typically used in retinotopic mapping studies. It might be interesting to repeat these experiments at increased spatial resolution in hopes of detecting existing topographic organization, if any.

Responses to squares at the ipsilateral horizontal meridian

As mentioned earlier, the mean ipsilateral versus ipsilateral predictions were better than chance. Conceivably, these results could have occurred given a scenario that would not reflect what is generally meant by ipsilateral selectivity. Namely, stimuli at the ipsilateral horizontal meridian might have elicited no or only negative activations, while ipsilateral stimuli near the vertical meridian might have elicited positive responses in neurons with large receptive fields centered contralaterally but extending across the vertical meridian. Therefore, we asked whether there actually were positive responses (as assessed by a GLM) to ipsilateral stimuli at the horizontal meridian. Simply averaging the GLM responses across the ROI to test whether the mean is positive would not be the best way to answer this question because true positive responses in some voxels could be washed out by negative responses in other voxels. Simply demonstrating that there are some positive GLM responses would not be the best way either because some positive responses could occur by chance. To circumvent these problems, we introduced an assumption—that there were true positive GLM responses to contralateral stimuli at the horizontal meridian—and tested the hypothesis that positive contralateral horizontal responses were greater than positive ipsilateral horizontal responses. A null result here would imply that there are positive GLM responses to stimuli at the ipsilateral horizontal meridian. To avoid selection bias in this analysis, responses from voxels with a positive response to contralateral squares were averaged for the contralateral mean, and responses from voxels with a positive response to ipsilateral squares were averaged for the ipsilateral mean. The two voxel sets were partially but not entirely overlapping. For this analysis, we used Experiment 3 data because no part of the square stimuli approached the vertical meridian. The activations used in the calculations were the outputs of a GLM in which regressors represented model responses to the entire series of stimulus presentations at each location. The results are shown in Figure 11. The only regions with a contralateral preference were the VC and the FG.

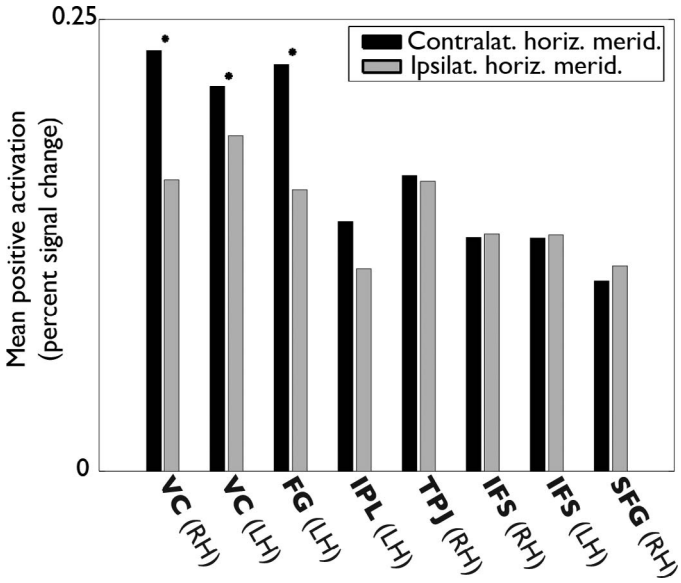

Figure 11. Activation levels in response to ipsilateral horizontal stimuli in Experiment 3. The height of each bar indicates the mean activity level of fMRI responses to contralateral and ipsilateral squares at the horizontal meridian. The means were calculated within ROI across subjects on positive fMRI activations only. That is, voxels with a positive response to contralateral squares were averaged for the contralateral mean, and voxels with a positive response to ipsilateral squares were averaged for the ipsilateral mean. Stars indicate a contralateral versus ipsilateral difference significantly ($p < 0.05$) greater than zero as measured by a one-tailed t test across voxels in all subjects. The activations used in the calculations were the outputs of a GLM in which regressors represented model responses to the entire series of stimulus presentations at each location. A contralateral preference was evident in the VC and FG only. In the remaining ROIs, positive activations to the ipsilateral horizontal meridian were roughly equal in magnitude to positive activations to the contralateral horizontal meridian.

Discussion

This study produced four main findings. First, visual stimuli, presented briefly and infrequently to fixating subjects, elicited fMRI activations in brain regions previously implicated in exogenous attention, namely the TPJ, IPL, and IFS. Second, MVPA showed that a classifier trained on data from these regions could distinguish between data points acquired when visual stimuli were presented in different parts of visual space. This finding demonstrates that the patterns of fMRI activity in the ROIs depended on the spatial location of the stimuli. Third, similar results were obtained using stimuli of different textures (expanding dots and radial checkerboard) and shapes (wedges, squares, and rings). These findings demonstrate that the spatial selectivity generalized across stimulus classes. Last, in the VC but not in the remaining ROIs, prediction accuracies were higher for spatially nonadjacent stimulus pairs than for adjacent stimulus pairs. The finding in the VC reflects the retinotopic organization known to exist there. The null results elsewhere indicate either that stimulus selectivity in these ROIs, as measured by our paradigm, is not organized retinotopically or that any existing organization is too coarse to be detectable by this method.

In the ROIs outside the VC, spatially selective voxels did not appear to be organized in retinotopic maps across the cortical surface. Instead, the sets of voxels active when stimuli appeared at different locations were distributed and overlapping, as in the spatial coding that has been described in the hippocampus and entorhinal cortex of the rat (O'Keefe, 2006). The distributed organization of spatial selectivity in the exogenous attention system, as measured with fMRI, is in contrast to the coarse retinotopy seen in the endogenous attention system (Silver & Kastner, 2009) and the relatively fine retinotopy seen in the early visual areas (DeYoe, Bandettini, Neitz, Miller, & Winans, 1994).

Among the regions where we documented selectivity for both contralateral and ipsilateral locations, organized in a distributed rather than topographic manner across the cortex, was the IFS. Our findings in the IFS differed from observations seen in previous studies, which documented contralateral selectivity organized as a rough topographic map across the region (Hagler & Sereno, 2006; Kastner & McCains, 2007). Note, however, that the stimuli and tasks in the earlier and current studies were very different. The earlier studies asked subjects to perform memory-guided saccades and working memory tasks using stimuli that were perfectly periodic and thus completely predictable. The current study used passive viewing of infrequently presented, unpredictable stimuli. Thus, the earlier studies implicitly encouraged the subjects to make use of goal-directed, endogenous attention, while the current study should have captured stimulus-driven, exogenous attention due to the stimulus unpredictability and rapid onset. We believe that this basic difference in experimental design is the likeliest explanation for the difference in results. We should note that the stimuli also evoked basic visual responses, and the stimuli conceivably may have caused subjects to endogenously or overtly attend to the locations of the stimuli given that the subjects were not performing a task at fixation. However, if endogenous attention or eye movements were consistently elicited by the stimuli, we would expect to see activations in brain regions implicated in endogenous attention and eye movements, namely the frontal eye field and (in humans) the medial bank of the IPS (Schluppeck et al., 2005; Sereno et al., 2001; Silver et al., 2005). Since we did not observe activations in

these brain regions, we conclude that it is unlikely that our observations elsewhere in the brain were driven by uncontrolled endogenous attention or eye movements.

The activation of the IFS by both the prior paradigms and our paradigm is consistent with the idea that this region acts as a transfer point between the stimulus-driven and goal-directed attention networks, participating in either network according to task demands (Serences et al., 2005). We suggest that anatomical connectivity between the IFS and the goal-directed attention network is organized differently from that between the IFS and the stimulus-driven attention network, producing differences in measured spatial selectivity in the IFS when one network or the other is preferentially activated.

The ROIs emerging from this study, like those in other studies manipulating exogenous attention, overlap substantially with regions where lesions often lead to persistent and severe spatial neglect—in particular, the right TPJ, IFS, and IPL (Behrmann & Tipper, 1999; Corbetta & Shulman, 2011; Karnath, 1997). Therefore, we were interested in evaluating our findings in the light of the spatial neglect literature. Spatial neglect is the failure to report, respond, or orient to novel or meaningful stimuli presented to the side of space opposite a brain lesion when this failure cannot be attributed to either elemental sensory or motor defects (Heilman, 1979; Heilman, Watson, & Valenstein, 2002). Typically, persistent and severe neglect is caused by damage to the right hemisphere. One neuroanatomical model of spatial neglect (Kinsbourne, 1977, 1993) proposes that the disorder arises because spatial representation in the damaged brain regions is contralateral in the left hemisphere and both contralateral and ipsilateral in the right hemisphere. In contrast to this model, we did not find that spatial selectivity in the left-hemisphere ROIs was limited to contralateral space. In fact, our classifier successfully distinguished every ipsilateral wedge or square from every other ipsilateral wedge or square in every ROI outside the VC, including (in the left hemisphere) the FG, IPL, and IFS. What did appear to be lateralized, however, was the overall level of activity in the TPJ: The right TPJ emerged as an ROI, but the left TPJ did not. This observation is consistent with several previous observations of right-hemisphere laterality in exogenous attention (Arrington et al., 2000; Downar et al., 2000; Shulman et al., 2010). Thus, our findings support the argument of Corbetta and Shulman (2011) that hemispheric asymmetry occurs in spatial neglect not because of an asymmetry in spatial attention per se but because the exogenous attention system is right lateralized and interconnected with a bilateral endogenous attention system that represents space contralaterally.

Keywords: exogenous, stimulus driven, involuntary, attention, spatial

Acknowledgments

We thank Gang Chen for valuable suggestions. This work was supported by the National Institute of Mental Health Intramural Research Program (NCT01087281; Protocol 93-M-0170).

Commercial relationships: none.
Corresponding author: Leslie G. Ungerleider.
Email: ungerlel@mail.nih.gov.
Address: Laboratory of Brain and Cognition, National Institute of Mental Health, National Institutes of Health, Bethesda, MD, USA.

References

Amedi, A., Malach, R., & Pascual-Leone, A. (2005). Negative BOLD differentiates visual imagery and perception. *Neuron, 48*(5), 859–872.

Arrington, C. M., Carr, T. H., Mayer, A. R., & Rao, S. M. (2000). Neural mechanisms of visual attention: Object-based selection of a region in space. *Journal of Cognitive Neuroscience, 12,* 106–117.

Behrmann, M., & Tipper, S. P. (1999). Attention accesses multiple reference frames: Evidence from visual neglect. *Journal of Experimental Psychology: Human Perception and Performance, 25,* 83–101.

Bressler, D. W., & Silver, M. A. (2010). Spatial attention improves reliability of fMRI retinotopic mapping signals in occipital and parietal cortex. *NeuroImage, 53,* 526–533.

Carlson, T. A., Schrater, P., & He, S. (2003). Patterns of activity in the categorical representations of objects. *Journal of Cognitive Neuroscience, 15,* 704–717.

Chen, Y., Namburi, P., Elliott, L. T., Heinzle, J., Soon, C. S., Chee, M. W., & Haynes, J. D. (2011). Cortical surface-based searchlight decoding. *NeuroImage, 56,* 582–592.

Cichy, R. M., Chen, Y., & Haynes, J.-D. (2011). Encoding the identity and location of objects in human LOC. *NeuroImage, 54,* 2297–2307.

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

Corbetta, M., & Shulman, G. L. (2011). Spatial neglect

and attention networks. *Annual Review of Neuroscience, 34,* 569–599.

Cox, D. D., & Savoy, R. L. (2003). Functional magnetic resonance imaging (fMRI) "brain reading": Detecting and classifying distributed patterns of fMRI activity in human visual cortex. *NeuroImage, 19*(2 Pt. 1), 261–270.

Cox, R. W. (1996). AFNI: software for analysis and visualization of functional magnetic resonance neuroimages. *Computers and Biomedical Research, 29*(3), 162–173.

Cox, R. W., & Hyde, J. S. (1997). Software tools for analysis and visualization of fMRI data. *NMR in Biomedicine, 10*(4–5), 171–178.

Dale, A. M., Fischl, B., & Sereno, M. I. (1999). Cortical surface-based analysis I: Segmentation and surface reconstruction. *NeuroImage, 9,* 179–194.

DeYoe, E. A., Bandettini, P., Neitz, J., Miller, D., & Winans, P. (1994). Functional magnetic resonance imaging (FMRI) of the human brain. *Journal of Neuroscience Methods, 54,* 171–187.

Dodds, C. M., Morein-Zamir, S., & Robbins, T. W. (2011). Dissociating inhibition, attention, and response control in the frontoparietal network using functional magnetic resonance imaging. *Cerebral Cortex, 21,* 1155–1165.

Downar, J., Crawley, A. P., Mikulis, D. J., & Davis, K. D. (2000). A multimodal cortical network for the detection of changes in the sensory environment. *Nature Neuroscience, 3,* 277–283.

Fischl, B., Sereno, M. I., & Dale, A. M. (1999). Cortical surface-based analysis II: Inflation, flattening, and a surface-based coordinate system. *NeuroImage, 9,* 195–207.

Fischl, B., Sereno, M. I., Tootell, R. B. H., & Dale, A. M. (1999). High-resolution intersubject averaging and a coordinate system for the cortical surface. *Human Brain Mapping, 8,* 272–284.

Hagler, D. J., Jr., Riecke, L., & Sereno, M. I. (2007). Parietal and superior frontal visuospatial maps activated by pointing and saccades. *NeuroImage, 35,* 1562–1577.

Hagler, D. J., Jr., & Sereno, M. I. (2006). Spatial maps in frontal and prefrontal cortex. *NeuroImage, 29,* 567–577.

Hansen, K. A., Kay, K. N., & Gallant, J. L. (2007). Topographic organization in and near human visual area V4. *Journal of Neuroscience, 27,* 11896–11911.

Hanson, S. J., Matsuka, T., & Haxby, J. V. (2004). Combinatorial codes in ventral temporal lobe for object recognition: Haxby (2001) revisited: Is there a "face" area? *NeuroImage, 23,* 156–166.

Hassabis, D., Chu, C., Rees, G., Weiskopf, N., Molyneux, P. D., & Maguire, E. A. (2009). Decoding neuronal ensembles in the human hippocampus. *Current Biology, 19,* 546–554.

Haxby, J. V., Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L., & Pietrini, P. (2001). Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science, 293,* 2425–2430.

Haxby, J. V., Horwitz, B., Ungerleider, L. G., Maisog, J. M., Pietrini, P., & Grady, C. L. (1994). The functional organization of human extrastriate cortex: a PET-rCBF study of selective attention to faces and location. *Journal of Neuroscience, 14*(11, Pt. 1), 6336–6353.

Haynes, J.-D., & Rees, G. (2006). Decoding mental states from brain activity in humans. *Nature Reviews Neuroscience, 7,* 523–534.

Heilman, K. M. (1979). Neglect and related disorders. In K. M. Heilman & E. Valenstein (Eds.), *Clinical neuropsychology* (pp. 268–307). New York, NY: Oxford University Press.

Heilman, K. M., Watson, R. T., & Valenstein, E. (2002). Spatial neglect. In H.-O. Karnath, D. Milner, & G. Vallar (Eds.), *The cognitive and neural bases of spatial neglect* (pp. 3–30). New York, NY: Oxford University Press.

Jimura, K., & Poldrack, R. A. (2012). Analyses of regional-average activation and multivoxel pattern information tell complementary stories. *Journal of Neuropsychologia, 50*(4), 544–552.

Jonides, J., & Yantis, S. (1988). Uniqueness of abrupt visual onset in capturing attention. *Perception & Psychophysics, 43,* 346–354.

Kamitani, Y., & Tong, F. (2005). Decoding the visual and subjective contents of the human brain. *Nature Neuroscience, 8,* 679–685.

Karnath, H. O. (1997). Spatial orientation and the representation of space with parietal lobe lesions. *Philosophical Transactions of the Royal Society of London B: Biological Sciences, 352,* 1411–1419.

Kastner, S., & McCains, S. A. (2007). Out of the spotlight: face to face with attention. *Nature Neuroscience, 10*(11), 1344–1345.

Kinsbourne, M. (1977). Hemi-neglect and hemisphere rivalry. *Advances in Neurology, 18,* 41–49.

Kinsbourne, M. (1993). Integrated cortical field model of consciousness. *Ciba Foundation Symposium, 174,* 43–50, discussion 51–60.

Konen, C., & Kastner, S. (2008). Representation of eye

movements and stimulus motion in topographically organized areas of human posterior parietal cortex. *Journal of Neuroscience, 28,* 8361–8375.

Kriegeskorte, N., & Bandettini, P. (2007). Analyzing for information, not activation, to exploit high-resolution fMRI. *NeuroImage, 38,* 649–662.

Levy, I., Schluppeck, D., Heeger, D. J., & Glimcher, P. W. (2007). Specificity of human cortical areas for reaches and saccades. *Journal of Neuroscience, 27,* 4687–4696.

Liu, T., Hospadaruk, L., Zhu, D. C., & Gardner, J. L. (2011). Feature-specific attentional priority signals in human cortex. *Journal of Neuroscience, 31,* 4484–4495.

Macaluso, E., Frith, C. D., & Driver, J. (2002). Supramodal effects of covert spatial orienting triggered by visual or tactile events. *Journal of Cognitive Neuroscience, 14,* 389–401.

Mur, M., Bandettini, P. A., & Kriegeskorte, N. (2009). Revealing representational content with pattern-information fMRI—An introductory guide. *Social Cognitive and Affective Neuroscience, 4,* 101–109.

Nichols, T. E., & Holmes, A. P. (2001). Nonparametric permutation tests for functional neuroimaging: A primer with examples. *Human Brain Mapping, 15,* 1–25.

Norman, K. A., Polyn, S. M., Detre, G. J., & Haxby, J. V. (2006). Beyond mind-reading: Multi-voxel pattern analysis of fMRI data. *Trends in Cognitive Science, 10,* 424–430.

O'Keefe, J. (2006). Hippocampal neurophysiology in the behaving animal. In P. Andersen, R. Morris, D. Amaral, T. Bliss, & J. O'Keefe (Eds.), *The hippocampus book* (pp. 475–548). New York, NY: Oxford University Press.

Schluppeck, D., Glimcher, P., & Heeger, D. J. (2005). Topographic organization for delayed saccades in human posterior parietal cortex. *Journal of Neurophysiology, 94,* 1372–1384.

Serences, J. T., Shomstein, S., Leber, A. B., Golay, X., Egeth, H. E., & Yantis, S. (2005). Coordination of voluntary and stimulus-driven attentional control in human cortex. *Psychological Science, 16,* 114–122.

Sereno, M. I., Pitzalis, S., & Martinez, A. (2001). Mapping of contralateral space in retinotopic coordinates by a parietal cortical area in humans. *Science, 294,* 1350–1354.

Shulman, G. L., Pope, D. L. W., Astafiev, S. V., McAvoy, M. P., Snyder, A. Z., & Corbetta, M. (2010). Right hemisphere dominance during spatial selective attention and target detection occurs outside the dorsal fronto-parietal network. *Journal of Neuroscience, 30,* 3640–3651.

Silver, M. A., & Kastner, S. (2009). Topographic maps in human frontal and parietal cortex. *Trends in Cognitive Sciences, 13,* 488–495.

Silver, M. A., Ress, D., & Heeger, D. J. (2005). Topographic maps of visual spatial attention in human parietal cortex. *Journal of Neurophysiology, 94,* 1358–1371.

Tootell, R. B., Silverman, M. S., Switkes, E., & De Valois, R. L. (1982). Deoxyglucose analysis of retinotopic organization in primate striate cortex. *Science, 218,* 902–904.

Tse, P. U. (2010). Attention underlies subjective temporal expansion. In A. C. Nobre & J. T. Coull (Eds.), *Attention and time* (pp. 137–150). New York, NY: Oxford University Press.

Vapnik, V. (1995). *The nature of statistical learning theory*. New York, NY: Springer.