

Spatially Global Representations in Human Primary Visual Cortex during Working Memory Maintenance

Edward F. Ester,¹ John T. Serences,² and Edward Awh¹

¹Department of Psychology, University of Oregon, Eugene, Oregon 97403, and ²Department of Psychology, University of California, San Diego, La Jolla, California 92037

Recent studies suggest that visual features are stored in working memory (WM) via sensory recruitment or sustained stimulus-specific patterns of activity in cortical regions that encode memoranda. One important question concerns the spatial extent of sensory recruitment. One possibility is that sensory recruitment is restricted to neurons that are retinotopically mapped to the positions occupied by the remembered items. Alternatively, specific feature values could be represented via a spatially global recruitment of neurons that encode the remembered feature, regardless of the retinotopic position of the remembered stimulus. Here, we evaluated these alternatives by requiring subjects to remember the orientation of a grating presented in the left or right visual field. Functional magnetic resonance imaging and multivoxel pattern analysis were then used to examine feature-specific activations in early visual regions during memory maintenance. Activation patterns that discriminated the remembered feature were found in regions of contralateral visual cortex that corresponded to the retinotopic position of the remembered item, as well as in ipsilateral regions that were not retinotopically mapped to the position of the stored stimulus. These results suggest that visual details are held in WM through a spatially global recruitment of early sensory cortex. This spatially global recruitment may enhance memory precision by facilitating robust population coding of the stored information.

Introduction

Working memory (WM) supports the maintenance of information in an online state. Human neuroimaging (Harrison and Tong, 2009; Serences et al., 2009a) and single-unit studies in nonhuman primates (Miller et al., 1996; Supér et al., 2001) suggest that WM storage is mediated by sensory recruitment or sustained activity in early cortical regions that encode information. For example, Serences et al. (2009a) required subjects to remember the orientation or color of a foveally presented grating. Using functional magnetic resonance imaging (fMRI) and multivoxel pattern analysis (MVPA), they demonstrated that sustained patterns of activation in primary visual cortex (V1) during WM storage discriminated the specific value of the remembered attribute, even in the absence of changes in overall response amplitude.

One important question concerns the spatial extent of sensory recruitment. On one hand, recruitment may be confined to cortical regions that were activated during the encoding of a stimulus. Dill and Fahle (1998) showed subjects two dot patterns separated by a short interval and asked them to determine whether these stimuli were identical. Relative to a condition in which the stimuli appeared in the same spatial location, perfor-

mance was reliably reduced when they appeared in different locations (Zaksas et al., 2001; Hollingworth, 2006, 2007). These findings suggest that object representations held in WM are spatially specific, thus motivating the hypothesis that sensory recruitment is confined to cortical regions that are retinotopically mapped to the spatial location occupied by the memoranda.

Alternatively, sensory recruitment may extend to cortical regions that were not engaged during stimulus encoding. For example, a specific orientation could be represented via the broad recruitment of cortical neurons that encode orientation, regardless of the retinotopic position of the remembered stimulus. This possibility is suggested by studies demonstrating a spatially global spread of feature-based attention during perception (Treue and Martinez-Trujillo, 1999; Serences and Boynton, 2007a). For example, Serences and Boynton (2007a) instructed subjects to monitor one of two directions of motion carried by stimuli located in one hemifield of a display. Using fMRI and MVPA, these authors demonstrated that patterns of activity in visual areas ipsilateral to the stimuli discriminated the attended direction of motion, despite the fact that these regions were not being driven by a bottom-up signal. During WM maintenance, this kind of spatially global recruitment might enhance the precision of mnemonic representations by recruiting additional neurons to facilitate robust population coding of the stored information (Pouget et al., 2003).

Here, we attempted to determine whether sensory recruitment during WM maintenance is spatially local or global. Subjects were required to remember the orientation of a grating presented in the left or right visual field. Using fMRI and MVPA, we found that patterns of activity in early cortical regions (V1)

Received Sept. 4, 2009; revised Oct. 10, 2009; accepted Oct. 26, 2009.

This work was supported by National Institutes of Health Grant R01MH087214 (E.A. and E.K.V.) E.A., J.T.S., and E.F.E. conceived and designed the experiment, E.F.E. collected and analyzed data, and E.F.E., J.T.S., and E.A. wrote the manuscript.

Correspondence should be addressed to Edward Awh, Department of Psychology, 1227 University of Oregon, Eugene, OR 97403. E-mail: awh@uoregon.edu.

DOI:10.1523/JNEUROSCI.4388-09.2009

Copyright © 2009 Society for Neuroscience 0270-6474/09/2915258-08\$15.00/0

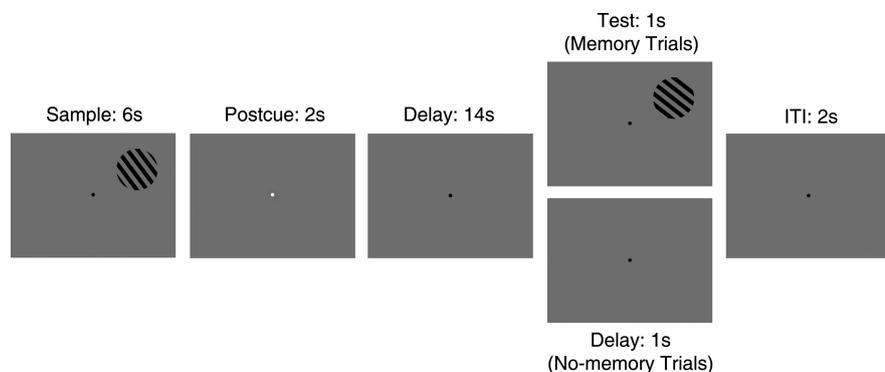


Figure 1. WM task. On each trial of the task, observers were shown a sample grating in the upper left or right visual field for 6 s. A subsequent change in the color of the fixation point (to red or green; shown here as white) informed observers whether they should remember the sample orientation (memory trials) or simply wait for the next trial to begin (no-memory trials). On memory trials, the cue was followed by a 14 s delay period and the presentation of a 1 s test grating in the same spatial location as the sample. Observers made a two-alternative, forced-choice response indicating whether the sample and test gratings shared the same orientation. On no-memory trials, the 14 s delay period was followed by an additional 1 s blank interval (i.e., no test grating was shown and no response was required). Trials were separated by a 2 s intertrial interval (ITI).

both contralateral and ipsilateral to a stimulus discriminated the remembered orientation. Under identical sensory conditions in which WM maintenance was not required, these effects were abolished. Furthermore, patterns of activity in contralateral and ipsilateral regions of interest (ROIs) were qualitatively similar during perception and WM maintenance, consistent with sensory recruitment accounts of WM storage.

Materials and Methods

Subjects. Twenty neurologically intact subjects participated in a single 2-h scanning session. All subjects gave informed written consent in accordance with the Office for the Protection of Human Subjects at the University of Oregon. Data from three subjects were discarded because of technical difficulties ($n = 2$) or voluntary withdrawal from the study ($n = 1$). Before scanning, each subject participated in a 1.5 h training session to identify performance thresholds in each behavioral task (see below, Staircase procedure). Subjects were compensated at a rate of \$8 per hour for behavioral testing and \$25 per hour for scanning.

Working memory task. Stimuli were generated in Matlab (version 7.1; MathWorks) using Psychophysics Toolbox software (Brainard, 1997; Pelli, 1997) and rendered in black on a medium-gray background via a rear projection system. Subjects were instructed to maintain fixation on a central dot (subtending 0.2° from a viewing distance of 58 cm) throughout each scan.

The sequence of events on each trial is presented in Figure 1. At the start of each trial, a “sample” grating stimulus (radius of 4° , 1 cycle/ $^\circ$) was displayed in the upper left or right portion of the display (horizontal and vertical eccentricity of $\pm 8.5^\circ$ and $+5^\circ$, respectively). The sample was rendered with an orientation of 45° or 135° (jittered on a trial-by-trial basis by a randomly selected value in the range of $\pm 10^\circ$ to discourage the use of categorical labels) and cycled on and off at a rate of 2 Hz (250 ms on, 250 ms off) for a total of 6 s. To attenuate the potency of retinal afterimages, the spatial phase of the sample was randomized on each cycle. After the sample epoch, a 2 s change in the color of the fixation point (to red or green) instructed subjects to remember the orientation of the sample for comparison with a subsequent test grating (a “memory” trial) or wait for the start of the next trial (a “no-memory” trial). For nine subjects, a green cue was used to denote memory trials, and a red cue denoted no-memory trials. For the remaining eight subjects, this mapping was reversed. The cue was followed by a 14 s delay period. On memory trials, the delay period was followed by the presentation of a static, 1 s test grating and a 2 s intertrial interval. Subjects made a button-press response during this 3 s interval to indicate whether the sample and test stimuli shared the same orientation (50% of trials). On no-memory trials, the delay period was instead followed by an additional 3 s of fixa-

tion. Stimulus orientation (45° or 135°), position (left or right visual field), and trial type (memory or no memory) were balanced within each scan (in which “scan” refers to a continuous, 404-s-long block of 16 trials). Subjects completed a total of six ($n = 1$), seven ($n = 7$), eight ($n = 7$), or nine ($n = 2$) scans as time permitted.

Perceptual monitoring task. During the same scanning session, subjects also performed a perceptual monitoring (PM) task. The stimuli and display parameters were closely modeled after those used in the WM task. On each trial, a sample grating oriented to 45° or 135° (randomly jittered on a trial-by-trial basis by a value in the range of $\pm 10^\circ$) was presented in the upper left or right visual field. This stimulus cycled on and off at a rate of 2 Hz (250 ms on, 250 ms off) for the duration of each 15 s trial. Subjects were instructed to monitor this stimulus and make a manual button-press response whenever they detected a brief (one cycle) change in its orientation (a target event; these occurred at unpredictable intervals two to

three times per trial; for information about the size of these angular deviations, see below, Staircase procedure). Each 260-s-long scan contained a total of 12 real and 3 “null” (i.e., 15 s of fixation) trials, and each observer completed a total of three ($n = 2$) or four ($n = 15$) scans as time permitted.

Staircase procedure. To ensure that both behavioral tasks were sufficiently challenging, we adjusted their difficulty for each subject during a separate behavioral testing session (completed 1–3 d before the scanning session). Sample-test (for the WM task) and sample-target (for the PM task) disparities were independently adjusted for each orientation category (45° or 135°) until a criterion level of performance (75% accuracy) was reached. The resulting orientation disparities were used to set sample-test and sample-target disparities during the scan session.

Eye tracking. To assess compliance with fixation instructions during PM and WM scans, eye position was visually monitored with the aid of an Applied Science Laboratories 5000 infrared tracking system in all subjects. For four subjects, eye position data during WM scans were explicitly recorded. This enabled more fine-grained assessments of whether there were any reliable differences in eye position as a function of the orientation stored in working memory. Recording was performed at 60 Hz, and data were filtered for blinks and corrected for linear drift offline.

fMRI data acquisition and analysis. fMRI data were collected using a 3 T Siemens Allegra system at the Robert and Beverly Lewis Center for Neuroimaging at the University of Oregon. Anatomical images were acquired using a spoiled-gradient-recalled T1-weighted sequence that yielded images with a 1 mm³ resolution. Whole-brain echo-planar images (EPIs) were acquired in 33 transverse slices (3 mm in-plane resolution, 2000 ms repetition time, 30 ms echo time, 90° flip angle, 64×64 matrix, 192 mm field of view, 3.5 mm slice thickness, no gap). EPIs were slice-time corrected, motion corrected (both within and between scans), and high-pass filtered (three cycles per run). Image preprocessing and data analysis were performed using BrainVoyagerQX (version 1.9) and custom time-series and pattern-classification routines written in Matlab.

Retinotopic mapping and voxel selection. Retinotopic mapping data were acquired using a rotating checkerboard stimulus flickering at 8 Hz and subtending 45° of polar angle (following Engel et al., 1994; Sereno et al., 1995). Each observer completed one scan lasting 480 s. This procedure was used to identify visual areas V1, V2v, V3v, V4v, V2d, V3d, and V3a in each cortical hemisphere. To aid in the visualization of these regions, data were projected onto a computationally inflated representation of each observer’s gray/white matter boundary.

To identify spatially selective voxels in these visual areas, we constructed a general linear model with a single boxcar regressor (denoting stimulus location, i.e., left vs right visual field) using data from all four PM scans. This regressor was then convolved with a gamma function to

account for the assumed shape of the hemodynamic response (Boynton et al., 1996). Voxels that showed a stronger response during epochs of contralateral (relative to ipsilateral) stimulation ($p < 0.05$, Bonferroni's corrected) were used to define ROIs in V1, V2v, V3v, hV4v, and V3a (as stimuli were presented in the upper visual field, ROIs in V2d and V3d typically contained too few voxels to enable a meaningful analysis) (for information on ROI sizes, see supplemental Table 1, available at www.jneurosci.org as supplemental material). The mean response amplitude across the 60 most spatially selective voxels in each visual region (collapsed across corresponding ROIs in each cortical hemisphere) identified using this method is shown in supplemental Figure 1 (available at www.jneurosci.org as supplemental material).

Multivoxel pattern analysis. MVPA provides a powerful complement to conventional univariate analyses of fMRI data. Posterior cortical regions such as V1 contain submillimeter columns of neurons that are selective for different stimulus features such as orientation. MVPA assumes that, if a particular fMRI voxel contains slight preponderance of columns tuned to a specific feature value, it should give rise to a weak but detectable response bias. By considering patterns across multiple weakly selective voxels, it is possible to infer the specific feature values that subjects are attending (Haxby et al., 2001; Haynes and Rees, 2005; Kamitani and Tong, 2005; Norman et al., 2006; Peelen and Downing, 2007; Serences and Boynton, 2007a,b; Serences et al., 2009b), imagining (Stokes et al., 2009), or remembering (Harrison and Tong, 2009; Serences et al., 2009a), even in the absence of sustained changes in overall response amplitude (Serences et al., 2009a).

To perform MVPA, we first extracted the raw time series from each voxel in a given ROI during a time period extending from 6 to 16 s after the offset of the sample stimulus (i.e., 12–22 s after the start of each trial). Each time series was normalized on a scan-by-scan basis using a z transform and sorted into one of eight bins according to three factors: ROI (contralateral or ipsilateral to the stimulus location), orientation (45° or 135°), and memory condition (memory or no memory). We then defined a "training" dataset using the data from all but one scan (here, scan refers to a single, continuous block of trials; thus, the training and test sets were always independent). Within the training dataset, we computed activation patterns comprising the mean response of each voxel during 45° and 135° trials. This was done independently for each location and memory condition pair, yielding a total of eight activation patterns. Data from the remaining scan were used to define a "test" dataset using an analogous approach. Before classification, we removed the overall mean from each activation pattern in the training and test sets. Finally, we trained a support vector machine (SVM) (specifically, the Ohio State University SVM implementation, <http://sourceforge.net/projects/svm>) to discriminate sample orientation using patterns from the training set and then used it to infer the orientation of the sample stimulus on each trial in the test set. Classification was performed separately for each stimulus location and memory condition pairing, so chance performance was always 50%. This analysis was iterated using a hold-one-out cross-validation procedure (Kamitani and Tong, 2005) until data from every scan had been used as the test set. Overall classification accuracy for each observer was then defined as the average classification accuracy for each observer across all six, seven, eight, or nine permutations (depending on how many WM scans the observer completed) of the hold-one-out procedure.

Results

The primary goal of this study was to use fMRI and MVPA to examine the spatial extent of sensory recruitment during WM maintenance. We examined response profiles within several regions of occipital cortex (V1, V2v, V3v, hV4v, and V3a); however, we focus on V1 because previous demonstrations show that this region displays robust patterns of orientation-selective activity during both perception (Haynes and Rees, 2005; Kamitani and Tong, 2005, 2006; Serences and Boynton, 2007a; Serences et al., 2009b) and WM maintenance (Harrison and Tong, 2009; Serences et al., 2009a). For all analyses, data are presented as collapsed across corresponding ROIs in each visual area (e.g., left

Table 1. Mean \pm SEM discrimination thresholds and accuracy for 45° and 135° trials during PM and WM maintenance

	PM task		WM task	
	45° trials	135° trials	45° trials	135° trials
Threshold	18.53 ± 1.49	18.40 ± 1.63	9.59 ± 1.26	9.89 ± 0.94
Accuracy	0.74 ± 0.01	0.72 ± 0.02	0.74 ± 0.02	0.74 ± 0.01

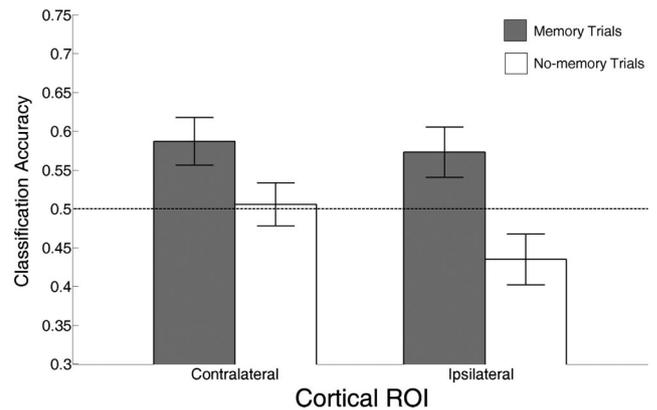


Figure 2. Multivoxel pattern analysis of WM data. The results of a multivoxel pattern analysis using the 60 most spatially selective voxels in each V1 ROI during a time period extending from 6 to 16 s after the offset of the sample stimulus. The horizontal line at 0.5 denotes chance classification accuracy. ROIs both contralateral and ipsilateral to the sample stimulus discriminated its orientation but only when active WM maintenance was required. Qualitatively similar results were obtained across a range of pattern sizes (50 and 75 voxels) (supplemental Fig. 2, available at www.jneurosci.org as supplemental material) and also when a linear discriminant algorithm was used to perform classification. Error bars represent ± 1 SEM.

and right V1) because no differences in classification accuracy were observed across cortical hemispheres. All analyses reported here are based on the 60 most spatially selective voxels in each ROI. Specific V1 ROIs in three subjects contained fewer than 60 voxels. For these subjects, analyses were performed using all available voxels within each ROI. Critically, we observed no systematic relationship between ROI sizes and overall classification accuracy, indicating that the findings reported here were not unfairly influenced by subjects with larger ROIs. In addition, all findings are robust across a broad range of pattern sizes (i.e., 50–75 voxels).

Behavioral performance in the PM and WM tasks is summarized in Table 1. Overall, orientation discrimination thresholds were higher during perceptual monitoring ($p < 0.001$). However, no differences in discrimination thresholds were observed between orientation categories (45° and 135°) in either task, and overall accuracy remained at a level near that dictated by the staircase procedure.

Multivoxel pattern analysis

If sensory recruitment operates in a spatially global manner, then sustained patterns of activation in ROIs both contralateral and ipsilateral to a remembered stimulus should discriminate its orientation. Figure 2 depicts the results of an MVPA based on the 60 most spatially selective voxels (as identified using data from the perceptual monitoring task; see Materials and Methods, Retinotopic mapping and voxel selection) from each V1 ROI during a time period extending from 6 to 16 s after the offset of the sample stimulus. An ANOVA revealed a main effect of memory condition ($F_{(1,16)} = 11.07$, $p < 0.001$), no effect of stimulus location (contralateral or ipsilateral; $F_{(1,16)} = 1.49$, $p > 0.24$), and no interaction between these two factors ($F_{(1,16)} < 1$). Planned com-

parisons revealed that classification accuracy was higher on memory trials relative to no memory trials. Critically, this was true in both contralateral ($t_{(16)} = 3.48, p < 0.01$) and ipsilateral ($t_{(16)} = 2.84, p < 0.05$) ROIs. Moreover, we failed to observe consistent above-chance classification in either ROI on no-memory trials ($t_{(16)} = 0.19, p > 0.8$ and $t_{(16)} = -1.99, p > 0.06$ for contralateral and ipsilateral ROIs, respectively) [apparent below-chance classification in ipsilateral ROIs during no-memory trials was not robust across variations in pattern sizes (supplemental Fig. 2, available at www.jneurosci.org as supplemental material)]. Because sensory conditions were equivalent during memory and no-memory trials, these findings suggest that the above-chance classification observed in both ROIs during memory trials reflects active WM maintenance rather than a lingering sensory response to the sample stimulus. Qualitatively similar findings were observed across a range of pattern sizes (50 and 75 voxels) (supplemental Fig. 2, available at www.jneurosci.org as supplemental material) and classification algorithms (linear discriminant function), indicating that above-chance classification is not idiosyncratically dependent on the use of an SVM.

MVPA was also applied to data from extrastriate visual areas V2v, V3v, hV4v, and V3a. However, contralateral and ipsilateral ROIs in each of these areas failed to support above-chance classification accuracy during memory or no-memory trials (supplemental Fig. 3, available at www.jneurosci.org as supplemental material) (for similar findings using foveal stimuli, see Serences et al., 2009a). Qualitatively similar results were observed over a wide range of pattern sizes (i.e., 50, 60, and 75 voxels) and different classification algorithms (e.g., linear discriminant function).

These results are consistent with the hypothesis that sensory recruitment during WM maintenance is spatially global. However, there are a number of possible issues that might also account for these effects. One such issue is head motion. Subtle differences in head position during 45° and 135° trials could differentially alter the intensity of fMRI voxels and bias the results of MVPA. To quantify any differences in head position as a function of the remembered orientation, we computed the mean change (relative to the start of each trial) in six motion parameters (translation and rotation in the x , y , and z planes) during a window extending from the 6–20s after the start of each memory trial (an earlier time window than that used for MVPA was selected to account for hemodynamic lag). These data were then submitted to a 2 (sample orientation, 45° or 135°) \times 2 (stimulus location, left or right visual field) \times 6 (motion parameter) ANOVA. This analysis revealed a main effect of visual field ($F_{(1,16)} = 17.49, p < 0.001$) and a marginal effect of motion parameter ($F_{(1,92,31,87)} = 2.24, p = 0.06$) (a Greenhouse–Geisser correction was applied to all model terms containing motion parameter data because of nonsphericity of this factor). Critically, this analysis failed to reveal any effects of stimulus angle ($F_{(1,16)} < 1$), nor did this factor interact with any other term (all $ps > 0.15$). As shown in supplemental Figure 4 (available at www.jneurosci.org as supplemental material), subjects exhibited a tendency to move or rotate in a direction consistent with the spatial location of a remembered stimulus. However, we failed to observe any biases in head position as a function of stimulus orientation, suggesting that above-chance orientation classification observed on memory trials was not solely attributable to subtle differences in head motion.

A second potential issue is eye position. Although subjects were instructed to maintain fixation for the duration of each WM scan (as assessed by visually monitoring ongoing eye position in all subjects with the aid of an Applied Science Laboratories track-

ing system), it is possible that they made subtle eye movements consistent with the orientation of a remembered grating. These eye movements could produce systematic modulations in activation patterns within retinotopically organized regions of visual cortex that could bias the results of an MVPA. To further assess compliance with fixation instructions and examine whether there were any subtle differences in eye position as a function of the remembered orientation, eye position data were recorded and tagged with trial-by-trial information regarding the specific angle stored for four subjects during scanning. To quantify differences in eye movements during WM scans, we identified all stable fixations (defined as a 200 ms epoch during which eye position did not deviate $>0.25^\circ$) that occurred outside of a 0.5° window encompassing the center of the screen during a period extending from 6 to 20s after the start of each memory trial. The x and y positions of fixations made by each of the four observers are shown in supplemental Figure 5 (available at www.jneurosci.org as supplemental material). Critically, we observed no differences in the mean or variance of these fixations as a function of remembered orientation (45° or 135°) across several parameters [x position, y position, polar distance (relative to the center of the screen), and polar angle; paired-samples t tests, all $ps > 0.3$], suggesting that above-chance classification accuracy during memory trials cannot be fully explained by subtle biases in eye movements.

Recent human neuroimaging (Serences and Boynton, 2007a) and single-unit (Martinez-Trujillo and Treue, 2004) studies in nonhuman primates have reported spatially global feature-based attention effects during perceptual processing. To examine whether similar effects could be observed in the present study, we performed a classification analysis on patterns observed in each ROI during a period 4–10s after the onset of the sample stimulus. Because the sample epoch of memory and no-memory trials were identical, both types of trial were included in this analysis. In contrast to the spatially global effects we observe during WM maintenance, neither contralateral nor ipsilateral V1 ROIs reliably discriminated the orientation of the sample stimulus (two-tailed t tests against chance, all $ps > 0.1$). We suspect that this lack of generality is attributable to the relatively low saliency of our stimuli. For example, MVPA was also applied to data from the PM task using a longer analysis window (6–16 s after stimulus onset). Following Serences and Boynton (2007a), patterns of activity in contralateral and ipsilateral V1 ROIs reliably discriminated the orientation of a monitored grating. However, when the same analysis was repeated using a shorter analysis window (4–10 s after stimulus onset), these effects were abolished. We speculate that, with a longer sample epoch, patterns of activity in both contralateral and ipsilateral ROIs would support robust classification of sample orientation.

Because classification was performed independently for each ROI, these findings do not establish whether contralateral and ipsilateral patterns within a given region are similar. To evaluate this possibility, we repeated our classification analysis using data from contralateral memory trials as the training set and examined whether this would allow accurate classification of the stored orientation when the same cortical ROI was ipsilateral to the sample stimulus (using the same hold-one-out cross-validation procedure described above; see Materials and Methods, Multivoxel pattern analysis). Across a range of pattern sizes (50, 60, and 75 voxels), this analysis failed to reveal above-chance classification accuracy (one-sample t tests against chance, all $ps > 0.1$). Thus, although patterns of activation in both contralateral and ipsilateral ROIs discriminate the orientation of a remembered

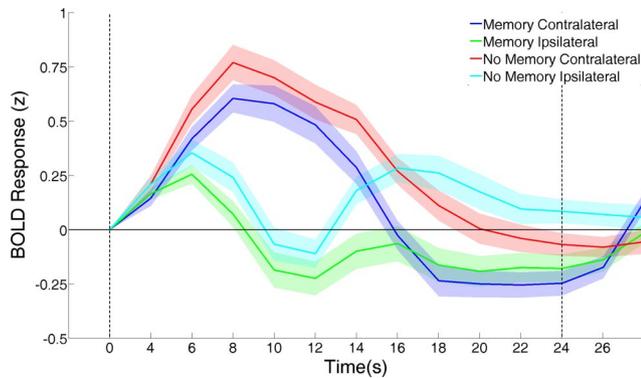


Figure 3. Event-related analysis of WM data. The mean response of the 60 most spatially selective voxels in V1 (collapsed across hemispheres) is shown as a function of memory condition and ROI. Data are collapsed across stimulus orientation (i.e., 45° vs 135° trials) because no differences in mean response amplitude attributable to this factor were observed. Dashed vertical lines at 0 and 22 s denote the onset of the sample and test stimuli, respectively. Shaded regions represent ± 1 SEM.

stimulus (Fig. 2), the pattern of activity in a given ROI is not necessarily identical across trials in which that region is contralateral and ipsilateral to the stored stimulus (for a qualitatively similar finding in the context of a perceptual discrimination task, see Serences and Boynton, 2007a). We speculate that this lack of generalization may be attributable to differences in lingering bottom-up activity when a given ROI is contralateral and ipsilateral to the stored stimulus. However, it is also possible that patterns of activation in ipsilateral ROIs may be qualitatively different from those observed in contralateral ROIs.

Event-related analysis of WM data

Data from the WM task were also submitted to a conventional univariate analysis. Figure 3 depicts the mean amplitude of the blood oxygenation level-dependent (BOLD) response in each V1 ROI during both memory and no-memory trials (no significant differences in mean response amplitude as a function of angle were observed; thus, the data are shown collapsed across this factor). To quantify differences between memory conditions and ROIs during WM maintenance, we computed the mean BOLD response within each condition during a period extending from 6 to 16 s after the offset of the sample stimulus (12–22 s in Fig. 2). A 2 (memory condition) \times 2 (ROI) ANOVA of these data revealed a main effect of ROI ($F_{(1,16)} = 11.69$, $p < 0.001$), a main effect of memory condition ($F_{(1,16)} = 8.45$, $p < 0.01$), and a significant interaction between these factors ($F_{(1,16)} = 10.72$, $p < 0.01$). Curiously, mean response amplitudes on no-memory trials were higher than those on memory trials (mean of 0.16 vs -0.02 , respectively, standardized units), particularly toward the end of the memory period. Mean response amplitudes were also greater in ipsilateral ROIs (mean of -0.08 vs -0.13 , for contralateral and ipsilateral ROIs, respectively), although this effect was not robust across variations in voxel number. Qualitatively similar effects were also observed in extrastriate visual areas (supplemental Fig. 7, available at www.jneurosci.org as supplemental material).

Similarities between spatially global signals during perception and WM maintenance

Recent studies (Harrison and Tong, 2009; Serences et al., 2009a) have reported a high degree of similarity between stimulus-specific patterns of activation in early sensory regions during perceptual processing and WM maintenance, consistent with the

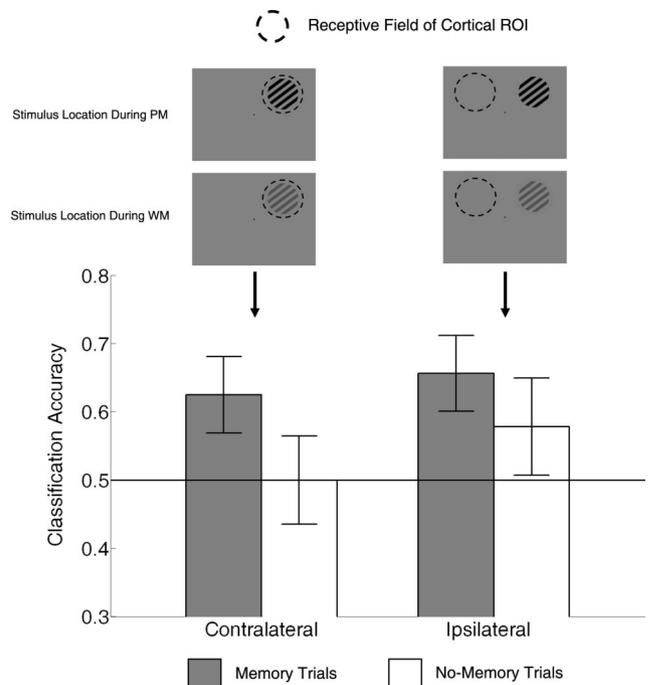


Figure 4. Global orientation-selective patterns are similar during PM and WM maintenance. A classification algorithm was trained to recognize orientation-selective patterns using the 75 most spatially selective voxels from each V1 ROI during PM (both left and right visual field trials were included in this analysis; top panels depict only RVF trials merely for illustrative purposes). As illustrated in the top two panels, this analysis was restricted to conditions in which the sample stimulus appeared in the same cortical ROI during PM and WM maintenance (for an identical analysis that focuses on conditions in which the sample stimulus appears in different spatial locations during each task, see supplemental Fig. 7, available at www.jneurosci.org as supplemental material). This algorithm was then used to infer the orientation of the sample stimulus during each trial of the WM task. The horizontal line at 0.5 denotes chance classification accuracy. Above-chance classification accuracy was observed in ROIs both contralateral and ipsilateral to the sample stimulus, indicating that orientation-selective patterns are similar during PM and WM maintenance. As in Figure 2, we failed to observe sustained above-chance classification accuracy on no-memory trials, indicating that above-chance classification of data from WM trials is not attributable to a lingering sensory effect. Qualitatively similar results were obtained when this analysis was applied to the 65 most spatially selective voxels in each ROI. Error bars represent ± 1 SEM.

hypothesis that WM involves the ingeneration of a specific perceptual event. For example, Serences et al. (2009a) trained a classification algorithm to recognize stimulus-specific patterns in V1 using data from a perceptually demanding orientation discrimination task and then used this classifier to successfully discriminate the identity of a remembered orientation during a separate WM task (for a similar demonstration, see Harrison and Tong, 2009). However, this analysis was restricted to cortical ROIs that corresponded to the location of the monitored or remembered stimulus. Here, we asked whether a qualitatively similar result would be observed in ROIs that were not retinotopically mapped to the position of the stimulus that was stored in WM (i.e., ROIs ipsilateral to the location of the stimulus). To examine this possibility, we trained a classifier to recognize stimulus orientation using data from PM scans. Recall that the stimuli used in PM scans were identical to those used in the WM task. However, during PM scans, stimuli cycled on and off at 2 Hz (250 ms on, 250 ms off) for the duration of each 15 s trial, and, thus, WM maintenance was not explicitly required. The classifier trained using the PM task was then used to infer the specific orientation that subjects were remembering during each trial of the WM task.

Figure 4 depicts the results of this analysis using the 75 most

spatially selective voxels from each cortical ROI. Data from one subject were excluded from this analysis as a result of 0% classification accuracy in each ROI across all tested pattern sizes. This was likely attributable to a large motion artifact between PM and WM scans that interfered with alignment of the functional images. The findings reported here represent the remaining 16 subjects. Collapsed across ROIs, this analysis revealed a successful generalization between patterns of activity observed during PM and WM maintenance but only when active memory maintenance was required (two-tailed *t* tests against chance, $p < 0.01$ and $p = 0.44$ for memory and no-memory conditions, respectively; in addition, a one-tailed paired-samples *t* test revealed significantly higher classification accuracy during memory trials, $t_{(15)} = 1.93$, $p < 0.05$). Qualitatively similar results were also observed when 65 voxels were used to perform this analysis. Thus, the pattern of activity observed within a specific ROI, whether it was contralateral or ipsilateral to the sample stimulus, was similar in the PM and WM tasks. These findings lend additional support to the hypothesis that WM involves the reiteration of a specific perceptual event. The same analysis failed to reveal consistent above-chance classification in extrastriate regions V2v, V3v, hV4v, and V3a.

These findings reflect conditions in which the sample stimulus appeared in the same spatial location during the PM and WM tasks. We next asked whether training on a contralateral stimulus in the PM task would allow successful classification of ipsilateral stimuli during the WM task. These data are shown in supplemental Figure 7 (available at www.jneurosci.org as supplemental material). This analysis failed to reveal consistent above-chance classification accuracy in any of the visual areas that we examined (two-tailed *t* tests against chance, all $ps > 0.15$), suggesting that similarities between PM and WM maintenance may be restricted to situations in which a monitored and remembered stimulus occupy the same spatial location.

Discussion

An emerging perspective, informed by human neuroimaging and single-unit recordings in nonhuman primates, is that WM storage is mediated by sensory recruitment or sustained activity in posterior cortical regions that encode memoranda (Miller et al., 1996; Awh and Jonides, 2001; Supèr et al., 2001; Jonides et al., 2005; Pasternak and Greenlee, 2005; Postle, 2006; D'Esposito, 2007). Here, we attempted to determine the spatial extent of this recruitment. Positional specificity effects in WM (Dill and Fahle, 1998; Zaksas et al., 2001; Hollingworth 2006, 2007) suggest that sensory recruitment is spatially local. However, several studies have reported spatially global feature-based attention effects during perceptual processing (Treue and Martinez-Trujillo, 1999; Saenz et al., 2002; Martinez-Trujillo and Treue, 2004; Bichot et al., 2005; Serences and Boynton, 2007a). Given known similarities between mechanisms of visual attention and working memory (Desimone, 1996; Rainer et al., 1998; Awh and Jonides, 2001), these findings suggest that sensory recruitment may be spatially global. Here, subjects were required to remember the orientation of a grating presented in one visual hemifield. Using fMRI and MVPA, we found that patterns of activity in early visual areas (e.g., V1) both contralateral and ipsilateral to this stimulus discriminated its orientation (Fig. 2). These findings cannot be explained by subtle differences in head or eye movements. Moreover, under identical sensory conditions in which WM maintenance was not explicitly required, these effects were abolished. Thus, above-chance classification observed on memory trials was not caused by a sustained sensory response to the sam-

ple stimulus. Additionally, we found that patterns of activity in contralateral and ipsilateral ROIs were qualitatively similar during perceptual processing and WM maintenance (specifically, activity patterns in contralateral ROIs during PM were qualitatively similar to patterns observed in contralateral ROIs during WM maintenance, whereas activity patterns observed in ipsilateral ROIs during PM were qualitatively similar to patterns observed in ipsilateral ROIs during WM maintenance) (Fig. 4), suggesting that similar mechanisms may support both processes. Together, these results suggest that sensory recruitment during WM maintenance is spatially global.

Several recent studies have reported successful classification of stimulus orientation in extrastriate visual areas during sensory encoding (Kamitani and Tong, 2005) and WM maintenance (Harrison and Tong, 2009). Likewise, studies of feature-based attention have documented spatially global effects in extrastriate visual areas (Treue and Martinez-Trujillo, 1999; Serences and Boynton, 2007a). However, the spatially global WM effects reported here were confined to primary visual cortex. We suspect that this difference is attributable to the relatively low salience of our stimuli. For example, the large foveal stimuli used by Harrison and Tong (2009) activated several hundred voxels within extrastriate visual areas V2–hV4v. In contrast, our small, peripheral stimuli produced a relatively weak response in only a few dozen voxels within each extrastriate visual area we examined (supplemental Table 1, available at www.jneurosci.org as supplemental material). It is possible that, with a larger response, patterns of activity in each of these visual areas would support robust classification of the remembered orientation.

Our findings complement recent studies that have reported feature-specific patterns of activation in retinotopically organized visual areas during WM maintenance in the absence of sustained changes in overall response amplitude (Harrison and Tong, 2009; Serences et al., 2009a). In the current study, we observed a relative decrease in response amplitude during WM maintenance (relative to a condition in which WM maintenance was not required). One possibility is that WM maintenance involves the recruitment of only those neurons tuned to the remembered feature value, whereas neurons tuned to other values are suppressed (relatively speaking). This might lead to an overall decrease in response amplitude during WM storage. However, we emphasize that this account is purely speculative because changes in response amplitude have been inconsistent across different studies that have examined activity in primary visual cortex. For example, Serences et al. (2009a) failed to observe sustained increases in response amplitude during WM maintenance (Offen et al., 2009). In contrast, Harrison and Tong (2009) observed sustained increases in response amplitude during WM maintenance in a subset of their subjects and no sustained changes in others. These studies, in concert with the findings reported here, raise questions about the extent to which sustained increases in response amplitude are diagnostic of the involvement of an area in WM maintenance.

Our findings are consistent with a growing number of studies that have reported spatially global signals in human and monkey retinotopic cortex during visual perception (Zaksas and Pasternak, 2005; Serences and Boynton, 2007a; Williams et al., 2008). In one example, Williams et al. (2008) required subjects to compare the category membership of two stimuli presented in the periphery of a visual display. Using fMRI and a multivariate analysis, these authors identified category-specific information in foveal retinotopic cortex, despite the fact that all stimuli were presented in spatial locations corresponding to cortical regions outside of this

area. One difference between the task devised by Williams et al. and the WM task used in this study is that, in the former, all stimuli were present for the duration of each trial. However, we speculate that similar feedback mechanisms support the spatially global mnemonic representations observed in this study. One possibility is that feedback directly facilitates stimulus-specific patterns of activity in both cortical hemispheres. Alternatively, feedback signals might be directed to a contralateral ROI and spread to ipsilateral ROIs via hard-wired cross-hemispheric connections between similarly tuned neurons in each hemisphere (for example, recent work has described a network of callosal axons linking cortical regions that represent the same orientation and spatial location of visual stimuli in cat visual areas 17 and 18) (Rocheffort et al., 2009). By this account, similarly tuned neurons in each cortical hemisphere are connected in a mutually excitatory manner; the efficacy of these connections might be modified by WM maintenance, giving rise to stimulus-specific modulations in the ipsilateral ROI. However, we believe this possibility unlikely given that patterns of activity in a given ROI during contralateral and ipsilateral trials are not necessarily identical (see Results, Multivoxel pattern analysis). Finally, inhibitory connections between corresponding visual areas in each hemisphere may play a role in producing stimulus-specific activation patterns. This possibility is consistent with the lack of generalization across patterns of activation evoked by contralateral and ipsilateral memoranda in a given ROI.

An important question concerns the boundary conditions that govern the spatially global mnemonic representations observed in this experiment. Several studies have demonstrated that performance on a memory-limited change detection task is degraded when the sample and test stimuli can occupy different spatial locations (Dill and Fahle, 1998; Zaksas et al., 2001). For example, Zaksas et al. (2001, their Experiment 1) trained monkeys to perform a change detection task while varying the spatial separation of the sample and test stimuli. These conditions were blocked, so it was always possible to infer the location of the test stimulus given the location of the sample. On some trials, the experimenters presented a random dot mask during the delay period; this stimulus could appear at the location of the sample stimulus, the impending test stimulus, or elsewhere in the display. Task performance was significantly impaired by the presentation of this mask but only when it appeared at the location of the impending test stimulus. The selective interference produced by this stimulus suggests that memory for the random dot aperture was localized in the cortical regions that were retinotopically mapped to the position of the impending test stimulus. However, as Zaksas et al. note, the masking effects observed in this study may simply mean that the discrimination judgment required by the task was biased toward information stored in the cortical regions that processed the eventual test location, even if the relevant information was also stored in other cortical regions. It is therefore still possible that the relevant directional information was stored in a spatially global manner during this task.

Do the spatially global activation patterns observed in this study have functional consequences for WM maintenance? Posterior visual areas V1–hV4v are thought to represent information about feature properties such as orientation via population response profiles (Pouget et al., 2000, 2003; Ma et al., 2006). One possibility is that the spatially global representations observed in this study act to enhance the precision of mnemonic representations (Awh et al., 2007; Zhang and Luck, 2008; Barton et al., 2009). For example, spatially global sensory recruitment would increase the number of neurons dedicated to representing a re-

membered feature attribute, which may in turn improve the efficiency or signal-to-noise ratio of population responses and thus enhance the precision of mnemonic representations. Future research is needed to examine this possibility.

In summary, our findings demonstrate that sensory recruitment during active WM maintenance is spatially global. We propose that spatially global gain modulations such as those observed in this study may serve to enhance the precision of mnemonic representations by recruiting additional sensory neurons that are not directly driven by the stimulus.

References

- Awh E, Jonides J (2001) Overlapping mechanisms of attention and spatial working memory. *Trends Cogn Sci* 5:119–126.
- Awh E, Barton B, Vogel EK (2007) Visual working memory represents a fixed number of items, regardless of complexity. *Psychol Sci* 18:622–628.
- Barton B, Ester EF, Awh E (2009) Discrete resource allocation in visual working memory. *J Exp Psychol Hum Percept Perform* 35:1359–1367.
- Bichot NP, Rossi AF, Desimone R (2005) Parallel and serial neural mechanisms for visual search in macaque area V4. *Science* 308:529–534.
- Boynton GM, Engel SA, Glover GH, Heeger DJ (1996) Linear systems analysis of functional magnetic resonance imaging in human V1. *J Neurosci* 16:4207–4221.
- Brainard DH (1997) The psychophysics toolbox. *Spat Vis* 10:433–436.
- Desimone R (1996) Neural mechanisms for visual memory and their role in attention. *Proc Natl Acad Sci U S A* 93:13494–13499.
- D'Esposito M (2007) From cognitive to neural models of working memory. *Philos Trans R Soc Lond B Biol Sci* 362:761–772.
- Dill M, Fahle M (1998) Limited translation invariance of human visual pattern recognition. *Percept Psychophys* 60:65–81.
- Engel SA, Rumelhart DE, Wandell BA, Lee AT, Glover GH, Chichilnisky EJ, Shadlen MN (1994) fMRI of human visual cortex. *Nature* 369:525.
- Harrison SA, Tong F (2009) Decoding reveals the contents of visual working memory in early visual areas. *Nature* 458:632–635.
- Haxby JV, Gobbini MI, Furey ML, Ishai A, Schouten JL, Pietrini P (2001) Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science* 293:2425–2430.
- Haynes JD, Rees G (2005) Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nat Neurosci* 8:686–691.
- Hollingworth A (2006) Scene and position specificity in visual memory for objects. *J Exp Psychol Learn Mem Cogn* 32:58–69.
- Hollingworth A (2007) Object-position binding in visual memory for natural scenes and object arrays. *J Exp Psychol Hum Percept Perform* 33:31–47.
- Jonides J, Lacey SC, Nee DE (2005) Processes of working memory in mind and brain. *Curr Dir Psychol Sci* 14:2–5.
- Kamitani Y, Tong F (2005) Decoding the visual and subjective contents of the human brain. *Nat Neurosci* 8:679–685.
- Kamitani Y, Tong F (2006) Decoding seen and attended motion directions from activity in the human visual cortex. *Curr Biol* 16:1096–1102.
- Ma WJ, Beck JM, Latham PE, Pouget A (2006) Bayesian inference with probabilistic population codes. *Nat Neurosci* 9:1432–1438.
- Martinez-Trujillo JC, Treue S (2004) Feature-based attention increases the selectivity of population responses in primate visual cortex. *Curr Biol* 14:744–751.
- Miller EK, Erickson CA, Desimone R (1996) Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *J Neurosci* 16:5154–5167.
- Norman KA, Polyn SM, Detre GJ, Haxby JV (2006) Beyond mind-reading: multi-voxel pattern analysis of fMRI data. *Trends Cogn Sci* 10:424–430.
- Offen S, Schluppeck D, Heeger DJ (2009) The role of early visual cortex in visual short-term memory and visual attention. *Vision Res* 49:1352–1362.
- Pasternak T, Greenlee MW (2005) Working memory in primate sensory systems. *Nat Rev Neurosci* 6:97–107.
- Peelen MV, Downing PE (2007) Using multi-voxel pattern analysis of fMRI data to interpret overlapping functional activations. *Trends Cogn Sci* 11:4–5.
- Pelli DG (1997) The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spat Vis* 10:437–442.
- Postle BR (2006) Working memory as an emergent property of the mind and brain. *Neuroscience* 139:23–38.

- Pouget A, Dayan P, Zemel R (2000) Information processing with population codes. *Nat Rev Neurosci* 1:125–132.
- Pouget A, Dayan P, Zemel RS (2003) Inference and computation with population codes. *Annu Rev Neurosci* 26:381–410.
- Rainer G, Asaad WF, Miller EK (1998) Selective representation of relevant information by neurons in the primate prefrontal cortex. *Nature* 393:577–579.
- Rochefort NL, Buzás P, Quenech'du N, Koza A, Eysel UT, Milleret C, Kisvárdy ZF (2009) Functional selectivity of interhemispheric connections in cat visual cortex. *Cereb Cortex* 19:2451–2465.
- Saenz M, Buracas GT, Boynton GM (2002) Global effects of feature-based attention in human visual cortex. *Nat Neurosci* 5:631–632.
- Serences JT, Boynton GM (2007a) Feature-based attentional modulations in the absence of direct visual stimulation. *Neuron* 55:301–312.
- Serences JT, Boynton GM (2007b) The representation of behavioral choice for motion in human visual cortex. *J Neurosci* 27:12893–12899.
- Serences JT, Ester EF, Vogel EK, Awh E (2009a) Stimulus-specific delay activity in human primary visual cortex. *Psychol Sci* 20:207–214.
- Serences JT, Saproo S, Scolari M, Ho T, Muftuler LT (2009b) Estimating the influence of attention on population codes in human visual cortex using voxel-based tuning functions. *Neuroimage* 44:223–231.
- Sereno MI, Dale AM, Reppas JB, Kwong KK, Belliveau JW, Brady TJ, Rosen BR, Tootell RB (1995) Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Science* 268:889–893.
- Stokes M, Thompson R, Cusack R, Duncan J (2009) Top-down activation of shape-specific population codes in visual cortex during mental imagery. *J Neurosci* 29:1565–1572.
- Supèr H, Spekreijse H, Lamme VAF (2001) A neural correlate of working memory in the monkey primary visual cortex. *Science* 293:120–124.
- Treue S, Martinez-Trujillo JC (1999) Feature-based attention influences motion processing gain in macaque visual cortex. *Nature* 399:575–579.
- Zaksas D, Pasternak T (2005) Area MT neurons respond to visual motion distant from their receptive fields. *J Neurophysiol* 94:4156–4167.
- Zaksas D, Bisley JW, Pasternak T (2001) Motion information is spatially localized in a visual working-memory task. *J Neurophysiol* 86:912–921.
- Zhang W, Luck SJ (2008) Discrete fixed-resolution representations in visual working memory. *Nature* 453:233–235.