

Population Dynamics of Early Visual Cortex during Working Memory

Masih Rahmati¹, Golbarg T. Saber², and Clayton E. Curtis¹

Abstract

■ Although the content of working memory (WM) can be decoded from the spatial patterns of brain activity in early visual cortex, how populations encode WM representations remains unclear. Here, we address this limitation by using a model-based approach that reconstructs the feature encoded by population activity measured with fMRI. Using this approach, we could successfully reconstruct the locations of memory-guided saccade goals based on the pattern of activity in visual cortex during a memory delay. We could reconstruct the saccade goal even when we dissociated the visual stimulus from the saccade goal using a memory-guided antisaccade procedure. By comparing the spatiotemporal population dynamics, we find that the representations in visual cortex are

stable but can also evolve from a representation of a remembered visual stimulus to a prospective goal. Moreover, because the representation of the antisaccade goal cannot be the result of bottom-up visual stimulation, it must be evoked by top-down signals presumably originating from frontal and/or parietal cortex. Indeed, we find that trial-by-trial fluctuations in delay period activity in frontal and parietal cortex correlate with the precision with which our model reconstructed the maintained saccade goal based on the pattern of activity in visual cortex. Therefore, the population dynamics in visual cortex encode WM representations, and these representations can be sculpted by top-down signals from frontal and parietal cortex. ■

INTRODUCTION

Working memory (WM) enables the brain to maintain behaviorally relevant information over short periods of time while the source of the information is no longer available (Curtis & D'Esposito, 2003; Baddeley & Hitch, 1974). Much of the previous work investigating the neural mechanisms that support WM has focused on the role of the frontal cortex (Srimal & Curtis, 2008; Goldman-Rakic, 1994; Funahashi, Bruce, & Goldman-Rakic, 1989). However, the contents of WM (e.g., orientation, color, spatial frequency, motion) can be decoded from delay period activity in human visual cortex (Albers, Kok, Toni, Dijkerman, & de Lange, 2013; Riggall & Postle, 2012; Ester, Serences, & Awh, 2009; Harrison & Tong, 2009; Serences, Ester, Vogel, & Awh, 2009). Findings such as these support the sensory recruitment model of WM, whereby the same neural mechanisms that encode sensory information are also used to maintain WM representations (D'Esposito & Postle, 2015; Postle, 2006).

Because these decoding studies depend on the patterns of activity across hundreds of voxels, visual cortex appears to encode WM representations within large populations of neurons. Multivariate pattern analysis (MVPA) (Haxby et al., 2001) uses powerful supervised machine learning algorithms to classify distinct brain states

(Haynes & Rees, 2006; Norman, Polyn, Detre, & Haxby, 2006; Kamitani & Tong, 2005). Although MVPA has successfully demonstrated that there are different patterns of activation in visual cortex associated with different WM features, it is limited in constraining the number of inferences one can draw about why these patterns are different (Serences & Saproo, 2012; Freeman, Brouwer, Heeger, & Merriam, 2011; Naselaris, Kay, Nishimoto, & Gallant, 2011). In contrast to MVPA, which attempts to predict brain states given a pattern of activity, an inverted encoding model (IEM) uses a set of a priori assumptions that form a linking hypothesis about how these brain states give rise to different patterns (Brouwer & Heeger, 2009; Dumoulin & Wandell, 2008). For example, if one assumes that a voxel's BOLD response during a spatial WM delay period is proportional to the responses of a set of information channels that tile visual space, the mapping between these channels and the BOLD responses can be used to predict the location of the memoranda in the visual field (Sprague, Ester, & Serences, 2014, 2016; Ester, Sprague, & Serences, 2015; Sprague & Serences, 2013). IEMs can reconstruct a full feature space, visual space in the above example, that can then be used for decoding a maintained stimulus feature (e.g., location).

Here, we model the patterns of neural population dynamics in early visual cortex to test several hypotheses. First, we ask if we can reconstruct the locations of memory-guided

¹New York University, ²Medical University of South Carolina

saccade goals based on the pattern of activity in visual cortex during a memory delay. Second, we ask if we can reconstruct the saccade goal even when we dissociate the location of the visual stimulus from the saccade goal using a memory-guided antisaccade procedure. Third, we test if the population dynamics evolve from a retrospective visual code to a prospective goal code. Fourth, because the representation of the antisaccade goal cannot be the result of bottom-up visual stimulation, it must be evoked by top-down signals, presumably originating from frontal cortex (Moore & Armstrong, 2003). To test this, we ask if trial-by-trial fluctuations in delay period activity in frontal cortex correlate with the precision with which our model reconstructs the maintained saccade goal based on the pattern of activity in visual cortex.

METHODS

Participants

Using procedures approved by New York University, five healthy participants (three women, right-handed, 25–42 years old) gave informed consent and participated in the study. Each volunteer participated in six scanning sessions: one

to obtain three high-resolution anatomies, one to measure retinotopy, and four to measure responses during memory-guided prosaccade and antisaccade tasks. Portions of the data were previously published in Saber, Pestilli, and Curtis (2015).

Memory-guided Saccade Tasks

Participants performed two memory-guided saccade tasks (Figure 1). Participants were instructed to remember the location of a briefly presented visual stimulus and immediately plan a saccade to, in the case of memory-guided prosaccades, or away from the stimulus, in the case of memory-guided antisaccades. The stimulus, a high-contrast flickering visual checkerboard (1° diameter, spatial frequency 2 cpd, 8 Hz flicker), appeared at one of eight locations separated by 30° angles at 10° eccentricity, avoiding vertical and horizontal meridians. Each trial began with 1.25 sec of central fixation followed by a brief presentation of the visual stimulus. Participants maintained the planned prosaccade or antisaccade over a long and variable delay period (3–15 sec). Disappearance of the fixation along with a brief auditory tone prompted




Figure 1. Experimental task, behavioral performance, and schematic of IEM. (A) Trial sequence of memory-guided prosaccade and antisaccade tasks. Participants remembered the location of a briefly presented stimulus in the periphery over a long and variable delay period and then generated a memory-guided saccade either toward the stimulus (prosaccade) or away from the stimulus (antisaccade). The dashed square depicted during the delay epochs represent the saccade goal but were not visible during the task. (B) Behavioral results of both the memory-guided pro- and antisaccade task. Neither SRTs nor saccade errors (both primary, MGS, and final eye position, FEP) differed significantly between the tasks. Error bars show *SEM* across participants. (C) The IEM. Regression weights (*W*) were calculated from a training set of BOLD data (*B*) and corresponding hypothetical channel coefficients (*C*). These regression weights were used to calculate the contribution of each basis function in the final reconstruction. The weighted linear combination of all basis functions was used to reconstruct a representation of visual space.

the memory-guided saccade. Participants held their gaze until the correct location was presented as feedback and corrected any error in their gaze. Participants then returned to fixation for the intertrial interval (10.5–15 sec). Partial trials were also included to help break the dependency of the delay on the target. These trials were aborted after the visual stimulus, signaled by a change in the color of the fixation spot. In four sessions, each with eight 322-sec scans, participants performed 192 prosaccade and 192 antisaccade trials, blocked by condition.

Oculomotor Methods

Eye position was recorded in the scanner at 1000 Hz (EyeLink 1000, SR Research, Ontario, Canada). We computed error in degrees of angle of the primary saccade as well as the final error after any quick corrections were made before the target feedback. All trials were inspected for noncompliance. Trials in which participants made an eye movement to the wrong saccade target were rare (0.3% and 0.6% of all pro- and antisaccade trials, respectively). We excluded trials where gaze deviated from fixation greater than 2° during the delay period or immediately following visual target (4.6% and 2.7% of pro- and antisaccade trials, respectively).

MRI Acquisition

All MRI data were acquired on a 3-T Allegra head-only scanner (Siemens, Erlangen, Germany) using a head transmit coil (NM-011) and two surface receive coils (NOVA Medical, Wakefield, MA): (1) a four-channel phased array receive surface coil (NMSC-021) for retinotopic mapping and (2) a four-element phased parallel array (NMSC-011) for memory-guided saccade tasks. We acquired T2*-sensitive echo-planar images (repetition time = 1.5 sec, echo time = 30 msec, flip angle = 75°, 26 slices, 3 × 3 × 3 mm voxels, 192 × 192 mm field of view). Three T1-weighted MPRAGE scans were averaged and used for gray matter segmentation, cortical flattening, registration, and visualization for creating each ROI (see below). Functional scans were corrected for head motion and aligned across sessions, detrended and high-pass filtered with a cutoff frequency of 0.0167 Hz, and converted to percentage signal modulation.

Retinotopic Mapping

We defined retinotopic visual areas (V1, V2, V3, V3A/B) and the first intraparietal sulcus map (IPS0) using a standard phase encoded approach (Wandell, Dumoulin, & Brewer, 2007). Polar angle components of the retinotopic maps were estimated with a high-contrast flickering wedge that rotated (clockwise or counterclockwise) around central fixation (subtended 90° of polar angle). Radial components were estimated with a high-contrast flickering annulus expanding or contracting from fixation

(subtended 4.2° radius). Each scan consisted of 10.5 cycles with a period of 24 sec and lasted 252 sec. Each participant completed six to eight runs of the rotating wedge aperture and four runs of expanding/contracting annulus aperture. We used flattened cortical surface representations of each participant's occipital and parietal cortices to visualize amplitude, coherence, and phase maps. Visual area boundaries were drawn on retinotopic maps based on standard conventions (Wandell et al., 2007; Larsson & Heeger, 2006). We used voxels from these ROIs in our IEM described below.

Estimating Delay Period Activity

We used a voxel-wise generalized linear model (GLM) to estimate the responses of each voxel to the stimulus encoding, memory retention interval (i.e., delay period), and memory-guided saccade intervals of each trial (Srimal & Curtis, 2008). Each epoch was modeled by the convolution of a canonical model of the hemodynamic impulse response function with a square wave (box-car regressor) whose duration was equal to the duration of the corresponding interval, detrended and high-pass filtered. Importantly, we estimated beta coefficients for each of the three trial epochs (stimulus, delay, response) for every trial independently. Modeling every event in each trial guaranteed that the beta coefficient corresponding to each epoch estimated the magnitude of BOLD activity specific to each epoch. As described below, we used the beta coefficients from the stimulus encoding epoch of the prosaccade task for model training and the beta coefficients from the delay periods of the prosaccade and antisaccade tasks for model testing.

Inverted Encoding Model

To reconstruct the spatial representation of the retinotopic maps, we used an IEM (Sprague & Serences, 2013). The IEM assumes that the activity of each neuron is a linear combination of the outputs of a set of basis functions, sometimes called information channels, tuned to different stimulus values. Similarly, it assumes that the activity of each voxel, modeled as the GLM beta coefficients, is a weighted sum of the activity of all neurons in that voxel. As a result, the activity of each voxel can be equated as a linear combination of the output of these basis functions for a given stimulus. This can be shown mathematically as

$$B = W \cdot C \quad (1)$$

where B is a matrix that represents the GLM beta coefficients corresponding to an event of interest (i.e., stimulus encoding, delay period) for all voxels in a given ROI and for all trials. W is a matrix of regression weights, and C contains the coefficients of the basis functions. The spatial representation of visual space is a linear combination of

all basis functions, each determined by its corresponding coefficient in C . To calculate the channel coefficients in a given trial using Equation 1, we use regression weights, W , that link the outputs of the basis functions to each voxel's activity. Thus, we can rewrite Equation 1 as

$$\hat{W} = B_{\text{train}} C_{\text{train}}^T (C_{\text{train}} C_{\text{train}}^T)^{-1} \quad (2)$$

where B_{train} in Equation 2 is a subset of the data used for training. However, during training we do not have access to the real contribution of each basis function to the population activity, and thus, we need to use a set of hypothetical channel coefficients, C_{train} . For each trial, we convolve the basis function (2-D Gaussian filter) with a Dirac delta function representing the stimulus location in the visual field. This allows us to estimate the overlap between each basis function and that part of the visual field covered by the stimulus in that trial. Using the regression weights, \hat{W} , and the population activity, B_{test} , in a given trial, we can calculate each channel's coefficients as

$$C_{\text{test}} = (\hat{W}^T \hat{W})^{-1} \hat{W}^T B_{\text{test}} \quad (3)$$

To avoid circularity in our analysis, we cross-validated our model by using the GLM-derived beta coefficients corresponding to the stimulus encoding epoch of the prosaccade task to calculate the regression weights, \hat{W} . We then used these regression weights to estimate the channel coefficients, C_{test} in Equation 3, given the population activity, B_{test} , during the delay period in the prosaccade and both encoding and delay periods in the antisaccade task. For the stimulus encoding epoch of the prosaccade task, we used a leave-one-out method to train and test the model. For each trial in the data set, we took that trial out and used the rest of trials to calculate the regression weights, \hat{W} . We then used these regression weights to estimate the channel coefficients, C_{test} in Equation 3, given the population activity, B_{test} , corresponding to the excluded trial. Thus, adding all basis functions, each multiplied by its corresponding channel coefficient in C_{test} , gives us the reconstructed spatial representation of visual space in that trial. Figure 1C shows a schematic of how our model works.

Based on the spatial selectivity of voxels in each ROI in early visual cortex, we used 56 2-D Gaussian filters with equal FWHMs distributed uniformly around a circle corresponding to a 10° eccentricity (i.e., eccentricity for the visual targets and saccade goals) to construct a set of basis functions. The number and width of these basis functions were chosen in a way to both avoid rank deficiency (caused by too much overlap between channels) and obtain smooth reconstructions. To make sure that using basis functions restricted to the stimulated ring does not bias the reconstructed spatial representation, we also tested a set of basis functions tiling the whole visual field. Because we did not find any meaningful differences in the reconstructed visuospatial maps, here we only report

the results from the basis functions restricted to the actual eccentricity used in the experiment.

Statistical Analysis

To evaluate the differences in behavioral performance between pro- and antisaccade tasks, we measured three metrics for each trial: the primary and final saccade accuracies and the RT. For each metric, we shuffled the data, mixing up the prosaccade and antisaccade trials, and calculated a t statistic that compared the shuffled "prosaccade" and "antisaccade" metrics. We did this 200 times to create a null distribution of t statistics that we used to compare to the t statistics computed using the veridically labeled data.

To verify that the IEM can reconstruct the WM content uniformly well across different saccade target locations, we plotted the model's reconstructed locations versus the true target locations. Within each topographic map, we calculated the peak of the reconstructed heat map averaged across all trials for each saccade target location and plotted this against the true target locations. Then, we estimated the slope of lines fit to these points, where the closer the slope of a line was to one indicated a more precise readout across different locations. Moreover, to evaluate the model performance on a trial-by-trial basis, we defined the readout precision at each trial as $1 - \frac{\text{error}}{180}$, where "error" is the difference between the angular position of the peak in the reconstructed retinotopic map and the actual location in the corresponding trial. We then measured the skewness of the distribution of precisions. The precision varies between 0 (maximum error) and 1 (minimum error). Thus, the higher the trial-by-trial model performance, the more the distribution of precisions was skewed toward 1. To estimate significance, we compared the skewness of the actual data with permuted distribution of skewness values. We shuffled the order of the trials 1000 times and, during each iteration, recomputed the skewness to form the permuted distribution of skewness values. Moreover, we identified correctly decoded trials based on the distance between the IEM readout peak and the actual target location in each trial. A trial is counted as a correctly decoded trial if the decoded location is within 15° , 30° , or 45° , as three defined correctness bounds, of the true location in that trial. We then, through a permutation test, compared the proportion of correctly decoded trials from the original and randomized order of data in each ROI. To do so, we shuffled the trial orders 1000 times, for each participant and each ROI, and reconstructed the visual-spatial maps each time. Then, we ran a Student's t test between the proportion of correctly decoded trials corresponding to the original order and shuffled.

We also tested the spatiotemporal interaction of the reconstructed representation of population activity under different conditions. To measure how reconstruction

precision changes from encoding to delay periods at the stimulus and saccade target locations, we fit a 2-D Gaussian to the reconstructed visual–spatial maps and measured the gain, width, and precision of the fit. We then used a two-way ANOVA to test the significance of the interaction between each of these parameters at two time epochs, the encoding and delay periods, and two locations, the stimulus location and saccade goal, which is the same on prosaccade but different on antisaccade trials.

Finally, to identify voxel clusters across the whole brain that significantly interact with model performance in an ROI, we calculated Pearson correlation between the delay period GLM beta coefficients, across trials for every voxel in the brain, and the trial-by-trial success of our model decoding of memorized space. We then used a permutation test to compare the z -transformation of above correlations corresponding to the original and randomized order of trials (corrected for multiple comparison between five ROIs). Next, we compared these clusters between all participants to find brain areas whose delay period activity correlates with the model precision in each of the five ROIs. We identified two clusters of voxels from two participants as the same if (1) the distance between the boundaries of two clusters was not larger than two voxels in the Talarach TT_N27 brain space and (2) both clusters were in the same brain area according to Freesurfer's cortical labels (*aparc.a2009s.annot*) from the automatic cortical parcellation.

RESULTS

Performance Equated on Memory-guided Pro- and Antisaccades

Figure 1B depicts the behavioral results for both pro- and antisaccade trials. Overall, 85% and 91% of all memory-guided prosaccade and antisaccade RTs (SRTs), respectively, fell within the normal range of visually guided saccades (i.e., 200–500 msec), and SRT did not significantly differ between the conditions (permutation test, $p = .17$). The accuracy of memory-guided prosaccades and antisaccades was also indistinguishable (permutation test, $p = .56$). Moreover, the final eye position after small corrective saccades (Mackey, Devinsky, Doyle, Golfinos, & Curtis, 2016; Mackey, Devinsky, Doyle, Meager, & Curtis, 2016), but before feedback, were not significantly different between the two conditions (permutation test, $p = .53$). We excluded trials from further analyses when the position of gaze deviated by more than 2° of visual angle away from the central fixation point during delay period. To test if smaller deviations from fixation are informative about the planned saccade location, we calculated the correlation between the angle of gaze deviations during different time epochs of the delay period and the angle of the saccade target. In none of the participants did these correlations reach significance (mean $r = -.05$, range: $r = -.09$ to $r = .10$). Thus, we were able to compare neural

differences unconfounded by behavioral differences between the prosaccade and antisaccade conditions and unconfounded by the influences of small gaze deviations.

Modeled Population Activity in Visual Cortex Reconstructs a Representation of WM

We collected BOLD fMRI signal while the participant performed memory-guided pro- and antisaccade tasks (Figure 1A). In both tasks, each trial can be divided into four epochs: central fixation, stimulation (stimulus encoding), delay period, and response. We mainly focused on stimulus encoding and delay period, as these are the intervals that the WM content is being encoded and maintained. We used an IEM (Sprague & Serences, 2013), described in detail in Methods, to reconstruct the spatial representation of neural population activity within each retinotopic map during the encoding and delay periods independently for both prosaccade and antisaccade trials (Figure 1C).

In the memory-guided prosaccade task, the participant is instructed to make a saccade toward the stimulus location after a variable delay period. We hypothesized that the activity patterns of neural population ensembles encode and maintain the stimulus location. Utilizing an IEM, we reconstructed the spatial representation of population activity in the five identified retinotopic areas. To average over all possible stimulus locations, we rotated the estimated channel coefficients, aligning all trials to one location, 45° in the upper right quadrant. Figure 2 shows the reconstruction results for a sample participant (Figure 2A) and for the average of all participants (Figure 2B) during both encoding and delay epochs. The lines on the polar plot depict the relative channel coefficients across the circular space of the model. The heat maps in the background show the reconstructed spatial representation of the neural population activity, which is the linear combination of the information channels each multiplied by its corresponding coefficients. In all retinotopic areas we identified, the IEM model yielded a high-fidelity reconstruction of the stimulus location during both time periods. This suggests that the population activity in these retinotopically organized cortical areas encode and maintain the stimulus location.

Next, we quantified the accuracy and precision of the model in each retinotopic area by estimating the trial-by-trial performance (Figure 3). For each trial, we estimated the accuracy by comparing the peak of the reconstructed visual space within each retinotopic area with the true stimulus location, separately at the encoding and delay epochs. To do this over all stimulus locations, we regressed the decoded locations, averaged over all trials with the same stimulus location, against the true locations. In all five retinotopic areas, we found a tight coupling between reconstructed location and true location as demonstrated by regression slopes that were near 1 (range across ROIs [0.94, 1.05], with largest $p = 1.6 \times 10^{-6}$). The top




Figure 2. The profile of channel coefficients and reconstructed visual space during memory-guided prosaccades. (A) Reconstruction for a single sample participant. In each retinotopic area, the mean channel coefficients (black lines) are plotted with *SEM* error bars across trials in gray shadows. To average across locations, the channels are rotated such that the true location is always at 45° . Notice that, in each retinotopic area, the population response peaks at the true location of the prosaccade target during both stimulus encoding and maintenance epochs and falls off as the angular distance from the true location grows. The heat maps in the background depicts the reconstructed spatial representation of the neural population activity, which is the linear combination of the information channels each multiplied by its corresponding coefficients. Again, for each trial, all targets are rotated to 45° . Notice that the peak of the heat map aligns with the true location of the prosaccade target. (B) The same as in A but averaged over all trials of all participants.




Figure 3. Model performance in decoding the true location of the stimulus during memory-guided prosaccades. (A) Performance from a single sample participant. The top row plots the mean ($\pm SEM$ across trials) decoded location of the stimulus against the true target location based on the model of V3 population activity during the stimulus encoding and delay epochs. The middle row plots the distribution of V3 model precision, where perfect decoding has a precision equal to 1. The bottom row plots the proportion of correctly decoded trials in each retinotopic area for different correctness boundaries (error range of 15°, 30°, or 45° of the true location). We consider a decoding correct if the readout error—the difference between the angular position of the decoded location and the true location—is smaller than a given correctness boundary. Overall, decoding was significantly above chance in all retinotopic areas during both the encoding and delay epochs. (B) The same as in A but averaged over all trials of all participants.

row of Figure 3 depicts the relation between the decoded location and true location in a sample retinotopic area (V3) during both encoding and delay epochs. We defined the precision of the model as the inverse of the decoded errors across trials, where smaller decoding error approached a precision of 1. The distribution of precision over trials was highly skewed toward 1 for each of the retinotopic areas (range across ROIs $[-1.36, -0.51]$, with largest $p = 1.2 \times 10^{-13}$; note that negative numbers indicate skewness toward 1). The middle row of Figure 3 depicts the distribution of precision in a sample retinotopic area (V3) during the encoding and delay epochs. Moreover, using various thresholds we asked how often did the model correctly decoded the stimulus location within a correctness boundary of 15°, 30°, or 45° of the true location. The bottom row of Figure 3 depicts the percentage of correctly decoded trials at different correctness boundaries. For all retinotopic areas, the percentage of correctly decoded trials was significantly above chance, even at the lowest interval of 15° (largest $p = 2 \times 10^{-118}$ for all correctness boundaries). Overall, our model of the population activity in these retinotopically organized cortical areas accurately and precisely represented the location of the encoded and maintained stimulus location.

Dissociating WM Representations of Past Visual Stimulation from Future Saccade Goals

In the memory-guided prosaccade task, the participant was instructed to plan and make a saccade toward the stimulus location. Therefore, we do not know if the population response encodes the visual target or the saccade goal because they are the same. To dissociate these two possible causes, we used a memory-guided antisaccade task in which the saccade goal is dissociated from the stimulus location (Figure 1A). We used the same IEM as in the prosaccade task to reconstruct the spatial representation of neural populations in each of the retinotopic areas during encoding and delay time epochs. To avoid any circularity in the reconstruction, we used the data corresponding to the encoding phase of the prosaccade task to calculate regression weights for the reconstruction at encoding and delay periods of the antisaccade task.

Figure 4 depicts the reconstructed retinotopic maps for both stimulus encoding and delay epochs for a sample participant and the average over all participants. As with the prosaccade data, the lines on the polar plot depict the relative channel coefficients across the circular




Figure 4. The profile of channel coefficients and reconstructed visual space during memory-guided antisaccades. (A) Reconstruction for a single sample participant. In each retinotopic area, the mean channel coefficients (black lines) are plotted with *SEM* error bars across all trials in gray shadows. To average across locations, the channels are rotated such that the true stimulus location is always at 45° and the true antisaccade goal is always at 135° . Notice that in each retinotopic area the population response peaks at both at the true location of the visual stimulus and the antisaccade target during both stimulus encoding and maintenance epochs and falls off as the angular distance from the true location grows. The heat maps in the background depicts the reconstructed spatial representation of the neural population activity, which is the linear combination of the information channels each multiplied by its corresponding coefficients. Notice that the peak of the heat map aligns with the true location of the visual stimulus and the antisaccade target. Moreover, the peak of the representation changes from the location of the visual stimulus during the encoding phase to the location of the antisaccade goal during the delay period. (B) The same as in A but averaged over all trials of all participants.




Figure 5. Model performance in decoding the location of the stimulus and saccade goal during memory-guided antisaccades. (A) Performance from a single sample participant. The top row plots the mean ($\pm SEM$ across trials) decoded location of the stimulus against the true target location based on the model of V3 population activity during the delay epoch. The middle row plots the distribution of V3 model precision, where perfect decoding has a precision equal to 1. The bottom row plots the proportion of correctly decoded trials in each retinotopic area for different correctness boundaries (error range of 15°, 30°, or 45° of the true location). We consider a decoding correct if the readout error—the difference between the angular position of the decoded location and the true location—is smaller than a given correctness boundary. Overall, decoding of stimulus and saccade target location was significantly above chance in all retinotopic areas. (B) The same as in A but averaged over all participants.

space of the model. Here, we aligned all trials such that the location of the visual stimulus was at 45° and the goal of the antisaccade was at 135°. In all retinotopic areas we identified, the IEM model reconstruction yielded two peaks, one at the location of the stimulus and another at the location of the saccade goal. Remarkably, this suggests that the population activity in these retinotopically organized cortical areas encode and maintain locations that were never visually stimulated and, therefore, must be the result of top-down stimulation.

As with the prosaccade data, we quantified the accuracy and precision of the model in each retinotopic area by estimating the trial-by-trial performance (Figure 5). Focusing solely on the delay period epoch, we compared separately both the decoded versus true location of the encoded stimulus and the decoded versus true location of the antisaccade goal. As depicted in the top row of Figure 5 for sample area V3, the true locations of the stimulus and antisaccade target were well predicted by the reconstructed locations as evidenced by regression slopes that were near 1 (range across ROIs [0.93, 1.12], largest $p = 1.0 \times 10^{-4}$). In each of the retinotopic areas, except V2, the distribution of precision over trials was highly skewed toward 1 when decoding the stimulus (range across ROIs [-0.5, -0.24], with

largest $p = .003$) and antisaccade target (range across ROIs [-1.0, -0.4], with largest $p = .009$) locations. For V2, the average skewness was -0.4 ($p = .17$). In a sample retinotopic area (V3), the middle row of Figure 5 depicts the distribution of precision of decoding the stimulus and antisaccade target locations based on the pattern of activity in the delay period. Next, we calculated the percentage of correctly decoded stimulus locations and correctly decoded antisaccade target locations within a boundary of 15°, 30°, or 45° of the true location. The bottom row of Figure 5 depicts the percentage of correctly decoded trials at different boundaries. For all retinotopic areas, the percentage of correctly decoded stimulus and antisaccade target locations was significantly above chance even at the lowest interval of 15° (largest $p = 8.7 \times 10^{-56}$ for all boundaries). Overall, our model of the population activity in these retinotopically organized cortical areas accurately and precisely represented the location of both the past stimulus location and future antisaccade goal.

Population Dynamics Track the Evolution of the Shifting Prioritized Location

To perform the antisaccade task, one must first encode the location of the visual stimulus and then compute,




Figure 6. Evolution of gain of readout peak in reconstructed visual space from stimulus location to saccade target during memory-guided antisaccades. (A) For a sample participant, each data point ($\pm SEM$) represents the gain of a 2-D Gaussian fit to the locations of the stimulus and the antisaccade target within the reconstructed visual space during the encoding and delay period epochs. Notice that in each retinotopic map the gain of the stimulus representation decreases from the encoding to delay period epoch, whereas the gain of the antisaccade target location increases from the encoding to delay period epoch. (B) The same as A but averaged over all participants.

plan, and maintain the location of the antisaccade target. As the trial evolves, the prioritized location therefore shifts from the stimulus location to the saccade goal. Consistent with this idea and as can be seen in Figure 4, it appears as if the reconstruction of the visual stimulus is greater during the encoding epoch than the delay period epoch. It also appears as if the reconstruction of the anti-saccade target is greater during the delay than the encoding epoch. Building on this observation, we quantified the gain of the stimulus and saccade target locations in the reconstructed visual space during the encoding and delay period epochs. To this end, we fit 2-D Gaussian functions to the peak locations in the reconstructions at both the stimulus and antisaccade target locations. In Figure 6, we plot the average gain at the stimulus and target locations during the encoding and delay epochs for each retinotopic area. As can be seen, there is a strongly significant reconstruction gain interaction between Epoch (encoding/delay) and Location (stimulus/saccade target) in all retinotopic areas (two-way ANOVA for V1, V2, V3, V3AB, and IPS0: smallest $F(1, 99) = 597, p = 10^{-81}$). This pattern supports the hypothesis that the spatiotemporal pattern of the population activity tracks the prioritized location as it evolves from the visual stimulus to the antisaccade goal.

Figure 7. Trial-by-trial fluctuations in delay period activity correlates with trial-by-trial fluctuations in the precision with which our model of population activity in early visual cortex could decode the maintained saccade. (A) For a single sample participant's right lateral hemisphere, the color overlay depicts a correlation map between the magnitude of delay period activity and model precision in decoding the memory-guided antisaccade target in area V3. Four clusters are highlighted with circles and nearby color-keyed histograms depict the distribution of correlation coefficients. Note that significant correlation coefficients, determined by permutation tests, are darkened. (B) The color overlay depicts the degree of overlap after averaging the correlation maps across all five participants. Although there was some variability across participants, there were portions of all four clusters where four out of five participants showed significant correlations with V3 decoding.




Table 1. Overlapping, between Four out of Five Participants, Clusters of Voxels Whose Activity Correlates with the Model Precision during Delay Period of the Antisaccade Task in Areas V2 and V3

<i>Hemisphere</i>	<i>Label</i>	<i>x</i>	<i>y</i>	<i>z</i>	<i>Volume (Voxels)</i>
<i>Clusters with Activity Correlating with Model Precision in Area V3</i>					
Right	G_occipital_middle	-36	77	-1	614
Right	G_front_inf-Opercular	-45	-11	8	568
Right	S_parieto_occipital	20	70	27	565
Left	S_temporal_sup	32	66	9	445
Right	Lat_Fis-post	-32	30	21	307
Right	S_central	-31	24	43	278
Right	S_postcentral	-17	47	48	220
Right	G_and_S_cingul-Ant	-6	-28	28	188
Right	G_occipital_middle	-47	74	2	176
Right	S_front_sup	-17	-5	46	176
Right	S_intrapariet_and_P_trans	-31	42	46	156
Right	S_front_sup	-16	-41	38	138
Left	G_front_sup	2	5	61	118
Right	G_front_middle	-30	37	57	108
Right	G_occipital_middle	-36	77	-1	614
<i>Clusters with Activity Correlating with Model Precision in Area V2</i>					
Right	S_central	-34	24	47	627
Right	S_calcarine	-23	57	3	500
Right	S_circular_insulasup;	-35	-6	9	450
Right	S_calcarine	-19	35	0	336
Right	G_front_middle	-35	-34	22	398
Left	S_circular_insula_sup	28	-11	12	236
Left	G_oc-temp_lat-fusifor	25	66	-15	226
Right	G_and_S_cingul-Ant	-19	-22	19	222
Right	S_postcentral	-27	38	39	218
Left	S_front_inf	41	-26	24	210
Left	G_and_S_cingul-Ant	13	-27	20	158
Right	S_oc-temp_lat	-39	52	-13	149
Right	S_central	-34	24	47	627

Coordinates are in the Talairach TT-N27 space, and labels are based on FreeSurfer *aparc.a2009s.annot* annotations.

Delay Period Activity in Frontal and Parietal Cortex Sculpt the Population Activity in Early Retinotopic Areas Improving the Precision of Readout

We demonstrated that the pattern of population activity in visual cortex represents the antisaccade goal. This representation is not the result of bottom-up visual stimulation and, therefore, must be instantiated by

top-down feedback signals. To identify the source of these top-down signals, we searched for brain areas whose trial-by-trial fluctuations in the magnitude of delay period activity correlate with the precision with which our model could accurately reconstruct the remembered location in the antisaccade task. We first, for each retinotopic area, created a vector of values that represented the model accuracy, inversely proportional

to the error of the reconstructed WM location. Second, using a GLM, we estimated, for all voxels in the brain, the delay period activity for each trial. Then, in an exploratory analysis, we calculated the correlation between decoding accuracy and delay period activity of each voxel. In individual participants, we found clusters of voxels whose delay period activity significantly predicted the decoding accuracy of the antisaccade goal from the pattern of activity in area V3 (Figure 7A). In the sample participant shown, we find significant clusters in dorsolateral pFC, superior precentral sulcus (SPCS), dorsal parietal cortex, and lateral occipital cortex. Despite some individual differences in the exact locations of significant voxel clusters within brain areas, we found consistent overlap in these four brain areas (Figure 7B). The histograms in Figure 7 depict the distribution of voxel correlations within each of the four brain areas, as well as their significance thresholds determined by permutation testing (Methods). In addition, Table 1 contains a list of all clusters of voxels in the brain, containing more than four contiguous voxels with overlap in at least four out of five participants, whose activity is significantly correlated with the decoding accuracy in areas V2 and V3. We focus here on areas V2 and V3 because the correlations were particularly strong and consistent, compared with the correlations with decoding from areas V1, V3AB, and IPS0, which was less consistent. The patterns of correlations identify both attentional and motor planning networks (Corbetta & Shulman, 2002) as potential sources for the modulatory top-down control signals prioritizing the saccade target location in the population activity in early visual cortex.

DISCUSSION

To summarize, we used model-based fMRI to reconstruct the locations of memory-guided saccade goals based on the pattern of activity in visual cortex during a memory delay. We could reconstruct the saccade goal even when we dissociated the location of the visual target from the saccade goal using a memory-guided antisaccade procedure. During these trials, the spatiotemporal population dynamics were informative, as they indicated an evolution from the past remembered visual stimulus to the future prospective goal. Early in the trial, the visual target location was primarily encoded in the population reconstruction. As the trial evolved in time, the population began to primarily encode the location of the saccade goal in the opposite hemifield. Therefore, the representation of the antisaccade goal cannot be the result of bottom-up visual stimulation and must be caused by top-down signals presumably originating from frontal and/or parietal cortex. In support of this hypothesis, we found that trial-by-trial fluctuations in delay period activity in frontal and parietal cortex predicts the precision with which our model reconstructed the maintained saccade goal based on the pattern of activity in visual cortex. Therefore, the pop-

ulation dynamics in visual cortex encode WM representations, and these representations can be sculpted by top-down signals from frontal and parietal cortex (Sreenivasan, Curtis, & D'Esposito, 2014; Curtis & D'Esposito, 2003).

Based on probabilistic population coding theory, we believe that neural populations do not encode sensory features—in our task the location of the saccade goal—but instead encode a probability distribution over sensory features (Ma, Beck, Latham, & Pouget, 2006; Zemel, Dayan, & Pouget, 1998). In this framework, the population characteristics, such as how the encoding of feature space is distributed across the topographic population and how a single neuron is tuned to variation over the feature space, determine the population dynamics. Using similar assumptions about the dynamics of populations of neurons, we built our IEM to decode spatial information encoded in the dynamics of populations of voxels in early visual cortex during WM. Using this model, we could successfully reconstruct a spatial representation of the saccade goal from the activity of neural ensembles in early visual cortex. The success of the model provides a strong linking hypothesis that connects decoding to population-level mechanisms. Specifically, we assumed an underlying neural architecture based on the retinotopic organization of the voxels within visual maps and modeled each voxel's response with a set of basis functions that tiled visual space. The success of spatial WM decoding depended on the fact that the precise structure of our model's basis set matched the representation of the feature space distributed across the neural population. Computational models of microcircuits describe how WM representations are maintained by the balance between recurrent local excitation among similarly tuned neurons and long-range inhibition (Wang et al., 2013; Lundqvist, Compte, & Lansner, 2010; Wang, 1999). Two key predictions from these models are relevant. First, WM representations are encoded in the collective response of populations of neurons whose tuning varies along the stimulus dimension. Second, persistent activity in the subset of neurons tuned to the remembered feature maintains the representation over time. In human visual cortex, evidence exists that support both of these predictions. The current study and other human neuroimaging studies have been successful at reconstructing representations of the population response encoding WM features (Sprague et al., 2014, 2016; Ester et al., 2015; van Bergen, Ma, Pratte, & Jehee, 2015; Ester, Anderson, Serences, & Awh, 2013; Sprague & Serences, 2013). Moreover, neural activity persists in early visual cortex during WM delays in humans, measured with BOLD imaging (Saber et al., 2015; Geng, Ruff, & Driver, 2009), and monkeys, measured with electrophysiology (van Kerkoerle, Self, & Roelfsema, 2017; Super, Spekreijse, & Lamme, 2001). Although the computational models were originally devised to model microcircuits in pFC, these results together open the possibility that similar mechanisms may support WM in early visual cortex.

During the memory-guided antisaccade trials, when we dissociated the location of the visual target from the saccade goal, the spatiotemporal population dynamics first encoded the visual target then encoded the saccade goal. Therefore, the representation of the antisaccade goal in early visual cortex must be the result of top-down signals and not some lingering visual response. Moreover, we recently demonstrated that fMRI activity persists in the specific parts of retinotopic maps that represent the location of memory-guided saccade goals (Saber et al., 2015). Whether these persistent BOLD responses in visual cortex reflect increases in spiking or only affect subthreshold synaptic activity remains unclear, but evidence for both exists (van Kerkoerle et al., 2017; Steinmetz & Moore, 2014; Zaksas & Pasternak, 2006; Super et al., 2001). Feedback signals from FEF increase the response gain of neurons in monkey areas V4 and MT (Merrikhi et al., 2017). Other studies confirm that the effects of top-down feedback signals originating from the monkey FEF are only detectable in extrastriate regions in the presence of bottom-up signals in visual cortex (sometimes referred to as “gating”; Ekstrom, Roelfsema, Arsenault, Bonmassar, & Vanduffel, 2008; Moore & Fallah, 2004; Moore & Armstrong, 2003). Together, these previous results suggest that the effects of feedback on visual cortex causes a multiplicative gain enhancement of visual responses. In our study of human visual cortex using fMRI, we find that top-down signals alone, without concomitant visual stimulation, are sufficient to induce a spatially specific pattern of topographic activity in visual cortex. Overall, our results are consistent with the sensory-motor recruitment model of WM that posits that the same mechanisms that evolved to create internal representations of sensory events are also used to maintain WM representations (Postle, 2006; Pasternak & Greenlee, 2005; Theeuwes, Olivers, & Chizk, 2005; Curtis & D’Esposito, 2003) and add to a growing list of studies demonstrating that WM representations are encoded in the patterns of activity in sensory cortices (Sprague et al., 2014; Albers et al., 2013; Ester et al., 2009, 2013; Lee, Kravitz, & Baker, 2013; Sprague & Serences, 2013; Christophel, Hebart, & Haynes, 2012; Riggall & Postle, 2012; Harrison & Tong, 2009).

Finally, we tested the hypothesis that the source of the top-down signals that prioritize the visual and saccade goal locations originate in the frontal and/or parietal cortex. The trial-by-trial fluctuations in delay period activity in both frontal and parietal cortex correlated with the precision with which our model of visual cortex reconstructed the maintained saccade goal. Notably, delay period activity near the junction of the sPCS and the superior frontal sulcus, the location of the putative human homologue of monkey area FEF, predicted the precision of encoding of the WM representation in visual cortex. Feedback signals during WM originating in FEF cause spatially selective increases in the response gain of neurons and also expand and shift the receptive fields of V4 and MT neurons toward the memorized location

(Merrikhi et al., 2017). Previous studies established the importance of the human sPCS for spatial WM. During spatial WM delays, neural activity in the human sPCS persists (Jerde, Merriam, Riggall, Hedges, & Curtis, 2012; Tark & Curtis, 2009; Simal & Curtis, 2008; Schluppeck, Curtis, Glimcher, & Heeger, 2006; Curtis, Rao, & D’Esposito, 2004). Moreover, TMS (Mackey & Curtis, 2017) and cortical resections (Mackey, Devinsky, Doyle, Meager, et al., 2016) to sPCS impact the precision of memory-guided saccades. Given that the sPCS is necessary for spatial WM, activity persists during WM maintenance, and it is topographically organized (Mackey, Winawer, & Curtis, 2017; Kastner et al., 2007), we propose that it maintains representations of spatial priority (Sprague & Serences, 2013; Jerde et al., 2012). Furthermore, priority maps may be the basis for top-down signals that sculpt the population dynamics in visual cortex in favor of neurons with receptive fields that match the locus of priority, providing a mechanism by which association cortex and sensory cortex interact to support WM.

Acknowledgments

This study was supported by R01-EY016407 to C. E. C. We thank the NYU Center for Brain Imaging’s Pablo Velasco, Keith Sanzenbach, and Valerio Luccio for help with data collection; John Serences and Thomas Sprague for help with the inverted encoding model; and Kartik Sreenivasan for helpful comments on the manuscript.

Reprint requests should be sent to Clayton E. Curtis, New York University, 6 Washington Place 859, New York, NY 10003, or via e-mail: clayton.curtis@nyu.edu.

REFERENCES

- Albers, A. M., Kok, P., Toni, I., Dijkerman, H. C., & de Lange, F. P. (2013). Shared representations for working memory and mental imagery in early visual cortex. *Current Biology*, *23*, 1427–1431.
- Baddeley, A. D., & Hitch, G. (1974). Working memory. *Psychology of Learning and Motivation*, *8*, 47–89.
- Brouwer, G. J., & Heeger, D. J. (2009). Decoding and reconstructing color from responses in human visual cortex. *Journal of Neuroscience*, *29*, 13992–14003.
- Christophel, T. B., Hebart, M. N., & Haynes, J. D. (2012). Decoding the contents of visual short-term memory from human visual and parietal cortex. *Journal of Neuroscience*, *32*, 12983–12989.
- Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, *3*, 201–215.
- Curtis, C. E., & D’Esposito, M. (2003). Persistent activity in the prefrontal cortex during working memory. *Trends in Cognitive Sciences*, *7*, 415–423.
- Curtis, C. E., Rao, V. Y., & D’Esposito, M. (2004). Maintenance of spatial and motor codes during oculomotor delayed response tasks. *Journal of Neuroscience*, *24*, 3944–3952.
- D’Esposito, M., & Postle, B. R. (2015). The cognitive neuroscience of working memory. *Annual Review of Psychology*, *66*, 115–142.
- Dumoulin, S. O., & Wandell, B. A. (2008). Population receptive field estimates in human visual cortex. *Neuroimage*, *39*, 647–660.

- Ekstrom, L. B., Roelfsema, P. R., Arsenault, J. T., Bonmassar, G., & Vanduffel, W. (2008). Bottom-up dependent gating of frontal signals in early visual cortex. *Science*, *321*, 414–417.
- Ester, E. F., Anderson, D. E., Serences, J. T., & Awh, E. (2013). A neural measure of precision in visual working memory. *Journal of Cognitive Neuroscience*, *25*, 754–761.
- Ester, E. F., Serences, J. T., & Awh, E. (2009). Spatially global representations in human primary visual cortex during working memory maintenance. *Journal of Neuroscience*, *29*, 15258–15265.
- Ester, E. F., Sprague, T. C., & Serences, J. T. (2015). Parietal and frontal cortex encode stimulus-specific mnemonic representations during visual working memory. *Neuron*, *87*, 893–905.
- Freeman, J., Brouwer, G. J., Heeger, D. J., & Merriam, E. P. (2011). Orientation decoding depends on maps, not columns. *Journal of Neuroscience*, *31*, 4792–4804.
- Funahashi, S., Bruce, C. J., & Goldman-Rakic, P. S. (1989). Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *Journal of Neurophysiology*, *61*, 331–349.
- Geng, J. J., Ruff, C. C., & Driver, J. (2009). Saccades to a remembered location elicit spatially specific activation in human retinotopic visual cortex. *Journal of Cognitive Neuroscience*, *21*, 230–245.
- Goldman-Rakic, P. S. (1994). Working memory dysfunction in schizophrenia. *The Journal of Neuropsychiatry and Clinical Neurosciences*, *6*, 348–357.
- Harrison, S. A., & Tong, F. (2009). Decoding reveals the contents of visual working memory in early visual areas. *Nature*, *458*, 632–635.
- 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.
- Haynes, J. D., & Rees, G. (2006). Decoding mental states from brain activity in humans. *Nature Reviews Neuroscience*, *7*, 523–534.
- Jerde, T. A., Merriam, E. P., Riggall, A. C., Hedges, J. H., & Curtis, C. E. (2012). Prioritized maps of space in human frontoparietal cortex. *Journal of Neuroscience*, *32*, 17382–17390.
- Kamitani, Y., & Tong, F. (2005). Decoding the visual and subjective contents of the human brain. *Nature Neuroscience*, *8*, 679–685.
- Kastner, S., DeSimone, K., Konen, C. S., Szczepanski, S. M., Weiner, K. S., & Schneider, K. A. (2007). Topographic maps in human frontal cortex revealed in memory-guided saccade and spatial working-memory tasks. *Journal of Neurophysiology*, *97*, 3494–3507.
- Larsson, J., & Heeger, D. J. (2006). Two retinotopic visual areas in human lateral occipital cortex. *Journal of Neuroscience*, *26*, 13128–13142.
- Lee, S. H., Kravitz, D. J., & Baker, C. I. (2013). Goal-dependent dissociation of visual and prefrontal cortices during working memory. *Nature Neuroscience*, *16*, 997–999.
- Lundqvist, M., Compte, A., & Lansner, A. (2010). Bistable firing and population oscillations in a modular attractor memory network. *PLoS Computational Biology*, *6*, e1000803.
- Ma, W. J., Beck, J. M., Latham, P. E., & Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature Neuroscience*, *9*, 1432–1438.
- Mackey, W. E., & Curtis, C. E. (2017). Distinct contributions by frontal and parietal cortices support working memory. *Scientific Reports*, *7*, 6188.
- Mackey, W. E., Devinsky, O., Doyle, W. K., Golfinos, J. G., & Curtis, C. E. (2016). Human parietal cortex lesions impact the precision of spatial working memory. *Journal of Neurophysiology*, *116*, 1049–1054.
- Mackey, W. E., Devinsky, O., Doyle, W. K., Meager, M. R., & Curtis, C. E. (2016). Human dorsolateral prefrontal cortex is not necessary for spatial working memory. *Journal of Neuroscience*, *36*, 2847–2856.
- Mackey, W. E., Winawer, J., & Curtis, C. E. (2017). Visual field map clusters in human frontoparietal cortex. *Elife*, *6*, e22974.
- Merrikhi, Y., Clark, K., Albarran, E., Parsa, M., Zirnsak, M., Moore, T., et al. (2017). Spatial working memory alters the efficacy of input to visual cortex. *Nature Communications*, *8*, 15041.
- Moore, T., & Armstrong, K. M. (2003). Selective gating of visual signals by microstimulation of frontal cortex. *Nature*, *421*, 370–373.
- Moore, T., & Fallah, M. (2004). Microstimulation of the frontal eye field and its effects on covert spatial attention. *Journal of Neurophysiology*, *91*, 152–162.
- Naselaris, T., Kay, K. N., Nishimoto, S., & Gallant, J. L. (2011). Encoding and decoding in fMRI. *Neuroimage*, *56*, 400–410.
- 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 Sciences*, *10*, 424–430.
- Pasternak, T., & Greenlee, M. W. (2005). Working memory in primate sensory systems. *Nature Reviews Neuroscience*, *6*, 97–107.
- Postle, B. R. (2006). Working memory as an emergent property of the mind and brain. *Neuroscience*, *139*, 23–38.
- Riggall, A. C., & Postle, B. R. (2012). The relationship between working memory storage and elevated activity as measured with functional magnetic resonance imaging. *Journal of Neuroscience*, *32*, 12990–12998.
- Saber, G. T., Pestilli, F., & Curtis, C. E. (2015). Saccade planning evokes topographically specific activity in the dorsal and ventral streams. *Journal of Neuroscience*, *35*, 245–252.
- Schluppeck, D., Curtis, C. E., Glimcher, P. W., & Heeger, D. J. (2006). Sustained activity in topographic areas of human posterior parietal cortex during memory-guided saccades. *Journal of Neuroscience*, *26*, 5098–5108.
- Serences, J. T., Ester, E. F., Vogel, E. K., & Awh, E. (2009). Stimulus-specific delay activity in human primary visual cortex. *Psychological Science*, *20*, 207–214.
- Serences, J. T., & Saproo, S. (2012). Computational advances towards linking BOLD and behavior. *Neuropsychologia*, *50*, 435–446.
- Sprague, T. C., Ester, E. F., & Serences, J. T. (2014). Reconstructions of information in visual spatial working memory degrade with memory load. *Current Biology*, *24*, 2174–2180.
- Sprague, T. C., Ester, E. F., & Serences, J. T. (2016). Restoring latent visual working memory representations in human cortex. *Neuron*, *91*, 694–707.
- Sprague, T. C., & Serences, J. T. (2013). Attention modulates spatial priority maps in the human occipital, parietal and frontal cortices. *Nature Neuroscience*, *16*, 1879–1887.
- Sreenivasan, K. K., Curtis, C. E., & D'Esposito, M. (2014). Revisiting the role of persistent neural activity during working memory. *Trends in Cognitive Sciences*, *18*, 82–89.
- Srimal, R., & Curtis, C. E. (2008). Persistent neural activity during the maintenance of spatial position in working memory. *Neuroimage*, *39*, 455–468.
- Steinmetz, N. A., & Moore, T. (2014). Eye movement preparation modulates neuronal responses in area V4 when dissociated from attentional demands. *Neuron*, *83*, 496–506.
- Super, H., Spekreijse, H., & Lamme, V. A. (2001). A neural correlate of working memory in the monkey primary visual cortex. *Science*, *293*, 120–124.

- Tark, K. J., & Curtis, C. E. (2009). Persistent neural activity in the human frontal cortex when maintaining space that is off the map. *Nature Neuroscience*, *12*, 1463–1468.
- Theeuwes, J., Olivers, C. N., & Chizk, C. L. (2005). Remembering a location makes the eyes curve away. *Psychological Science*, *16*, 196–199.
- van Bergen, R. S., Ma, W. J., Pratte, M. S., & Jehee, J. F. (2015). Sensory uncertainty decoded from visual cortex predicts behavior. *Nature Neuroscience*, *18*, 1728–1730.
- van Kerkoerle, T., Self, M. W., & Roelfsema, P. R. (2017). Layer-specificity in the effects of attention and working memory on activity in primary visual cortex. *Nature Communications*, *8*, 13804.
- Wandell, B. A., Dumoulin, S. O., & Brewer, A. A. (2007). Visual field maps in human cortex. *Neuron*, *56*, 366–383.
- Wang, M., Yang, Y., Wang, C. J., Gamo, N. J., Jin, L. E., Mazer, J. A., et al. (2013). NMDA receptors subserve persistent neuronal firing during working memory in dorsolateral prefrontal cortex. *Neuron*, *77*, 736–749.
- Wang, X. J. (1999). Synaptic basis of cortical persistent activity: The importance of NMDA receptors to working memory. *Journal of Neuroscience*, *19*, 9587–9603.
- Zaksas, D., & Pasternak, T. (2006). Directional signals in the prefrontal cortex and in area MT during a working memory for visual motion task. *Journal of Neuroscience*, *26*, 11726–11742.
- Zemel, R. S., Dayan, P., & Pouget, A. (1998). Probabilistic interpretation of population codes. *Neural Computation*, *10*, 403–430.