Feedback scales the spatial tuning of cortical responses during visual memory

Robert Woodry[1], Clayton E. Curtis[1,2], Jonathan Winawer[1,2]

1 Department of Psychology, New York University, New York City, NY 10003

2 Center for Neural Science, New York University, New York City, NY 10003

Keywords: working memory, long-term memory, fMRI, retinotopy, saccades

Abbreviated title: *Cortical tuning of visual memory*

Number of pages: 33

Number of figures: 7

Number of tables: 0

Word Count

Abstract: 240

Introduction: 648

Discussion: 1494

Conflicts of interest: The authors declare no competing financial interests.

Acknowledgements: We thank New York University's Center for Brain Imaging for technical support. This research was supported by the National Eye Institute (R01 EY033925 & R01 EY016407 to C.E. Curtis; R01 EY027401 to J. Winawer; P30 EY013079 Core grant for vision) by the National Institute of Mental Health (R01 MH111417 to J. Winawer), and pilot funds from the NYU Center for Brain Imaging to R. Woodry. We thank Ilona Bloem, Jan Kurzawski, Marc Himmelberg, Rania Ezzo, and Ekin Tunçok for help with scanning; Serra Favila, Tommy Sprague, and Ekin Tunçok for helpful comments on the manuscript; and Serra Favila for sharing data and code. Because this is my first submission I'd like to extend special thanks to Kate Yurgil and Liz Chrastil for taking a chance on me.

Corresponding Author: Robert Woodry (rfw256@nyu.edu)

Abstract

Perception, working memory, and long-term memory each evoke neural responses in visual cortex, suggesting that memory uses encoding mechanisms shared with perception. While previous research has largely focused on how perception and memory are similar, we hypothesized that responses in visual cortex would differ depending on the origins of the inputs. Using fMRI, we quantified spatial tuning in visual cortex while participants (both sexes) viewed, maintained in working memory, or retrieved from long-term memory a peripheral target. In each of these conditions, BOLD responses were spatially tuned and were aligned with the target's polar angle in all measured visual field maps including V1. As expected given the increasing sizes of receptive fields, polar angle tuning during perception increased in width systematically up the visual hierarchy from V1 to V2, V3, hV4, and beyond. In stark contrast, the widths of tuned responses were broad across the visual hierarchy during working memory and long-term memory, matched to the widths in perception in later visual field maps but much broader in V1. This pattern is consistent with the idea that mnemonic responses in V1 stem from top-down sources. Moreover, these tuned responses when biased (clockwise or counterclockwise of target) predicted matched biases in memory, suggesting that the readout of maintained and reinstated mnemonic responses influences memory guided behavior. We conclude that feedback constrains spatial tuning during memory, where earlier visual maps inherit broader tuning from later maps thereby impacting the precision of memory.

Significance Statement

We demonstrate that visual information that is seen, maintained in working memory, and retrieved from long-term memory evokes responses that differ in spatial extent within visual cortex. These differences depend on the origins of the visual inputs. Feedforward visual inputs during perception evoke tuned responses in early visual areas that increase in size up the visual hierarchy. Feedback inputs associated with memory originate from later visual areas with larger receptive fields resulting in uniformly wide spatial tuning even in primary visual cortex. That trial-to-trial difficulty is reflected in the accuracy and precision of these representations suggests that visual cortex is flexibly used for processing visuospatial information, regardless of where that information originates.

Introduction

While it is clearly established that occipital cortex contains visual maps (Inouye, 1909), recent evidence supports the provocative idea that these visual maps also play important roles in both working and long-term memory. Surprisingly, the contents of visual working memory (Harrison and Tong, 2009; Serences et al., 2009; Curtis and Sprague, 2021) and the contents retrieved from long-term memory (Bosch et al., 2014; Naselaris et al., 2015; Vo et al., 2022) can be decoded from the patterns of voxel activity in human primary visual cortex (V1). Such results provide strong support for influential theories of how the encoding mechanisms used for perception might also be used to store working memory representations (D'Esposito and Postle, 2015; Serences, 2016) and similarly be used to recall the visual properties of retrieved long-term memory (Tulving and Thomson, 1973; Schacter et al., 1998; Rugg et al., 2008). Thus, it is no longer a question as to whether visual cortex participates in cognitive functions beyond perception, but a question of how.

Patterns of evoked activity during perception can be used to predict the contents of memory (Albers et al., 2013; Bosch et al., 2014; Rademaker et al., 2019), supporting the idea that they share encoding mechanisms in visual cortex. Moreover, recalling large objects evokes activity that encroaches into more peripheral portions of visual field maps (Kosslyn et al., 1995), as if recalled objects that are larger encompass more of the visual field just like it does for seen objects. Despite these seeming parallels, we hypothesized that these responses in striate and extrastriate cortex also differ because the origins of their inputs differ. During perception, visual information is transmitted through the eyes and the retinogeniculate pathway to a cluster of retinotopically organized maps, including V1, V2, and V3 (Van Essen and Maunsell, 1983). Neural activity in these early visual maps strongly influences one's percepts (Tong et al., 1998). During working memory, information enters visual cortex in the same feedforward way, but after stimulus offset the maintenance of visual information depends on interactions between higher order cortical regions like the prefrontal cortex and sensory areas like V1. Such interactions support the storage of working memory representations (Curtis and D'Esposito, 2003; D'Esposito and Postle, 2015; Curtis and Sprague, 2021). During long-term memory, information is thought to be reconstructed in visual cortex through the retrieval of a memory stored in other brain structures such as the hippocampus (Schacter et al., 1998).

Most of the research described above focused on similarities between memory and perceptual representations. Some recent work, however, has also observed systematic differences between visual cortex activations between perception and long-term memory (Breedlove et al., 2020; Favila et al., 2022), and between perception and working memory (Rademaker et al., 2019; Kwak and Curtis, 2022). But several important questions remain. There have been no comparisons of visual cortex responses in perception, working memory, and long-term memory with the same study parameters, limited links between the precision of long-term memory reactivation and single-trial behavior, and no comparisons of the temporal dynamics between working memory and long-term memory signals in visual cortex.

Here, we use the receptive field properties of neuron populations within voxels (Dumoulin and Wandell, 2008) in visual field maps to quantify and compare the spatial tuning of evoked responses during perception and memory. As a preview, fMRI responses during perception matched the spatial position of seen targets and increased in tuning width up the visual hierarchy consistent with increases in receptive field sizes (Smith et al., 2001; Dumoulin and Wandell, 2008). During working and long-term memory, tuning widths were large in early visual cortex as if they were inherited from feedback from higher order areas with large receptive fields. Critically, errors in these spatially tuned responses during visual memory aligned with errors in memory behavior, suggesting that memory behavior depends on a readout of these maintained and retrieved memory responses in early visual cortex.

Materials and Methods

Subjects

Eight human subjects (5 Males, 25-32 years old) were recruited to participate in the experiment and were compensated for their time. Subjects were recruited from the New York University community and included author R.F.W. Other subjects were naive to the purpose of the experiments. All subjects gave written informed consent to procedures approved by the New York University Institutional Review Board prior to participation. All subjects had normal or corrected-to-normal visual acuity, normal color vision, and no MRI contraindications. No subjects were excluded from the main data analyses.

Stimuli

Target stimuli

For each of the three conditions (perception, working memory, long-term memory), 16 polar angles were selected from 16 evenly spaced bins along an isoeccentric ring at 7° eccentricity from a central fixation point (0-22.5 deg, 22.5-45 deg, … 337.5 to 360 deg). Within each bin, the precise target location was randomly assigned. The 48 target locations were uniquely generated for each participant. The target stimulus at each spatial location consisted of a drifting Gabor patch (sd = 0.33 deg, spatial frequency: 2 cyc/deg, truncated at 5 sd). Each Gabor patch drifted radially towards fixation at a rate of 3 Hz (1.5 deg/s), with an orientation tangential to the isoeccentric ring.

Object stimuli

A small image of an object (1° diameter) was shown at fixation at the beginning of each trial. For the long-term memory condition, these pairings stayed consistent throughout the pre-scan behavioral training (see below) and during the main experiment, and therefore a particular object served as a cue for the location to retrieve from long-term memory. For perception and working memory, the pairings were unique for each trial, ensuring that subjects could not form associations between objects and target locations. The paired stimuli were selected randomly without replacement from a bank of colored images consisting of everyday objects (BOSS dataset, (Brodeur et al., 2010).

Experimental procedure

Main experiment

The main experiment cycled through three scans: one each for perception, working memory, and long-term memory, repeated twice per scan session (Figure 1). Each subject participated in two scan sessions spaced no more than a week apart, for a total of 12 scans (4 per condition). Each scan had 16 trials corresponding to the 16 target locations, in random order. Across both sessions, this culminated in a total of 64 trials per condition, 192 trials total per subject.

For perception trials, we instructed participants to fixate on a central cross. The cross was briefly replaced by an object for 0.5 seconds. After the object disappeared, the cross re-appeared and the target stimulus remained visible for 11.5 seconds, during which subjects maintained central fixation. After this delay period, the fixation cross changed from black to green, indicating to the subject to make a saccade to the target location. Participants were instructed to maintain their fixation at the expected location until the fixation cross changed color back to black, at which point they returned their gaze to the central fixation cross. This saccade response period lasted 1.5 seconds for each trial.

The working memory block was the same except the target stimulus disappeared when the object disappeared. The subject was instructed to "hold the target in mind" throughout the delay while centrally fixating. At the end of the delay, the subject was similarly cued to make a saccade response to the location of the target stimulus.

The long-term memory block was the same except that the target stimulus was not shown at all. The subject was instructed to retrieve from memory the target stimulus associated with the object, with these associations learned during a pre-scan behavioral training session (see next section). Similar to the working memory and perception blocks, participants maintained central fixation during the delay period of 11.5 seconds and made a saccade to the target location when the cross turned green.

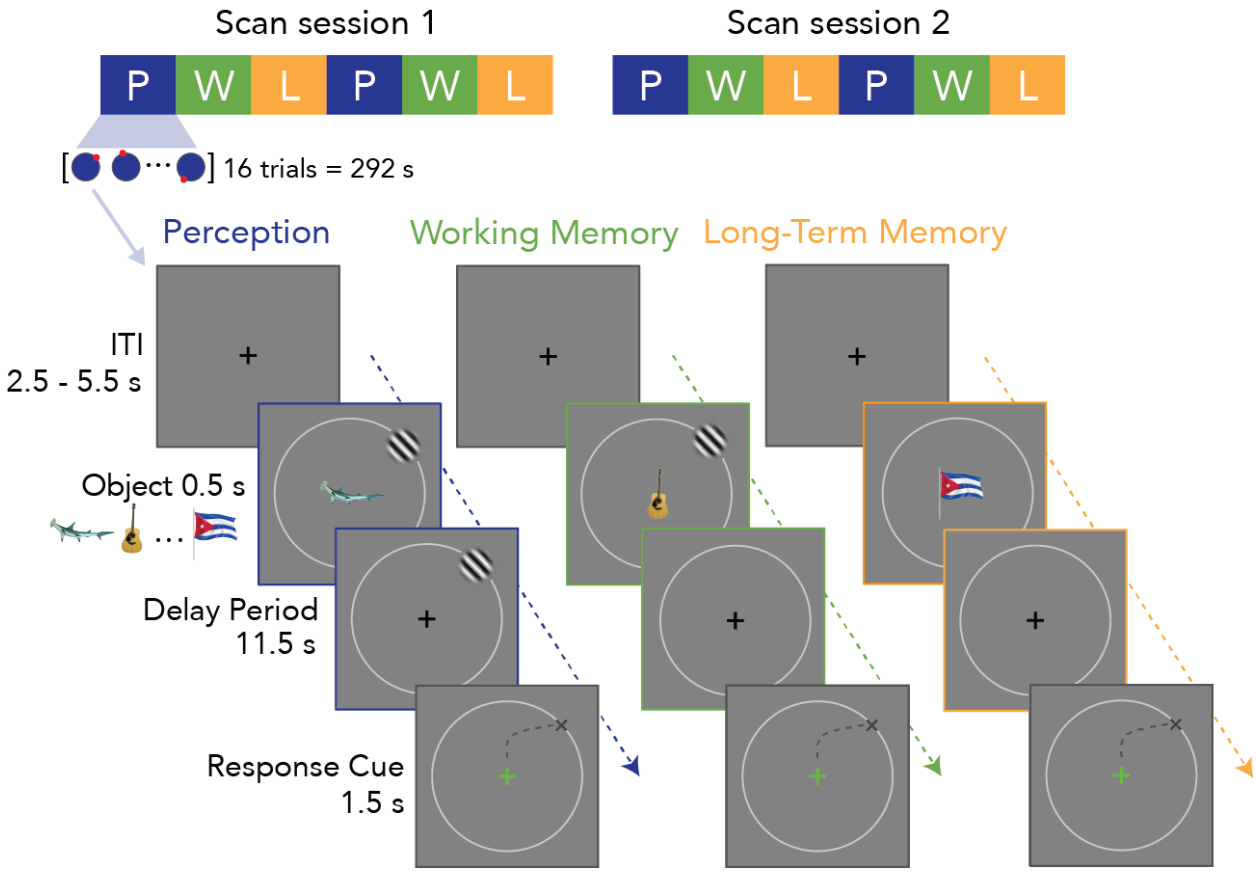

Figure 1. Main fMRI task design.

Subjects participated in two fMRI sessions with 6 scans each: perception, working memory, and long-term-memory scans, repeated twice (top). Each 292-s scan included 16 trials corresponding to 16 polar angles, in random order, always at 7° eccentricity. Objects are shown centrally for 0.5 seconds, but are only meaningful for the long-term memory trials. In Perception trials, the target remains on screen for the 11.5-s delay. For working memory trials, the target disappears after 0.5 s. For long-term memory trials, the target is not shown. At the end of the delay, the fixation cross turns green and participants make a saccade to the target location. The conditions are matched except for how the target location is accessed and how it is maintained throughout the delay period.

long-term memory training.

Subjects learned associations between 16 object stimuli and 16 target locations by completing study and retrieval blocks before each scan session (outside the scanner). Following each study block, subjects completed three retrieval blocks. In each trial of the study block, an object stimulus was briefly presented at fixation simultaneously with its corresponding Gabor target stimulus (Figure 2a). Subjects were instructed to fixate a central cross and learn the association between each object stimulus and its corresponding target location. The study block

was self-paced with a minimum 1 second inter-trial interval (ITI). Each of the 16 object/target pairs was presented five times per block (80 trials), with at least two study blocks per behavioral training session.

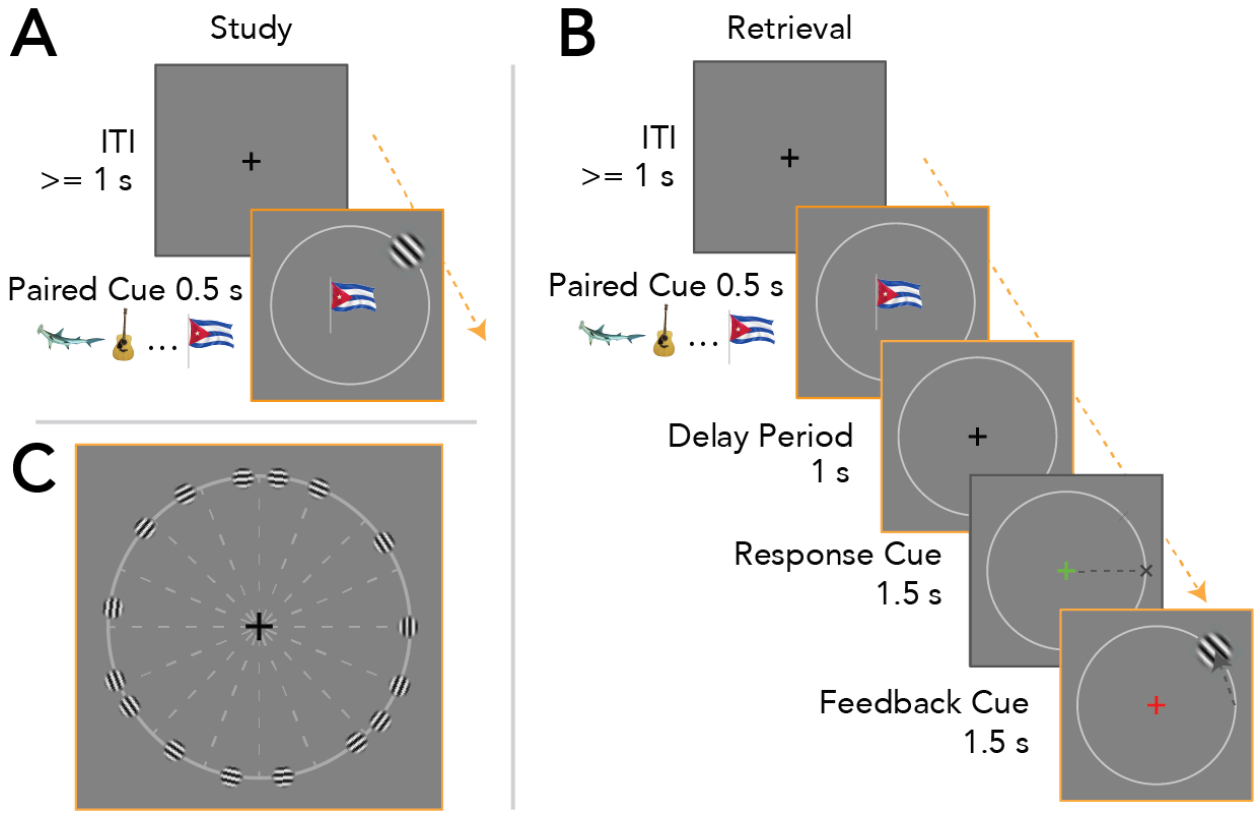

Figure 2. Prescan long-term memory training.

A) Study phase: Self-paced brief (0.5 second) viewings of the paired pre-cue at fixation and the associated target Gabor in the periphery, followed by an inter-trial interval (min. one second) Each participant had their own set of 16 unique pairs to memorize. B) Retrieval phase: The bulk of learning happened during this training phase. The retrieval phase was similarly self-paced, where following an inter-trial interval (min. one second), a paired cue was briefly presented at fixation without the target in the periphery. After a brief delay period the participant was cued to make a saccade response to the target's location. Feedback was then given in the form of the fixation cross either remaining green if the saccade was closest to the target's location, or changing to red if the response landed closer to a different target's location. C) Example layout of target stimuli: For each condition, target stimuli are presented at 7° eccentricity, sampled around the visual field within each of 16 non-overlapping 22.5° bins. Every participant had their own unique set of target locations. This led to some targets spaced near each other (i.e. near the bin borders marked by the dashed lines), while others were spaced far apart.

In the retrieval block, subjects were instructed to fixate the central cross while the object stimulus was flashed briefly, followed by a short delay where the subject was instructed to

retrieve from memory the associated target stimulus (Figure 2b). This was followed by the fixation cross changing color to green, which indicated to the participant to make an eye movement to the target's expected location. Feedback followed each retrieval trial in the form of the fixation cross remaining green if correct or changing color to red if incorrect. A saccade was considered correct if it was closer to the correct target than to any of the other 15 possible targets, otherwise it was considered incorrect. Simultaneously, the target stimulus was revealed at its true location, with the subject instructed to make a corrective saccade to the target stimulus. The retrieval block was also self-paced with a 1 second minimum ITI. Each of the 16 object/target pairs was presented four times per block (64 trials), with at least six retrieval blocks for the first training session, and four for the second session. Subjects were asked after six blocks if they felt they learned all the associations or if they needed to do more practice. If they needed more practice, they did at least one more retrieval block.

Retinotopic mapping procedure

In addition to the main experiment each participant completed 10-12 retinotopic mapping scans in a single separate scan session. The stimuli and procedures are the same used by Himmelberg et al. (2021), described here in brief. Each scan consisted of contrast patterns windowed by bar apertures (1.5 deg width) that swept across the visual field within a 12-deg-radius circle. There were 8 sweeps along different directions. While vertical and horizontal sweeps traveled the entire extent of the circular aperture, the diagonal sweeps stopped halfway, and were then replaced by blank periods. Each bar sweep took 24 s to complete. At 8 sweeps for each functional run, each scan took 192 s in total to complete. The contrast patterns were pink noise (grayscale) background with randomly placed and sized items. The stimuli and background were updated at 3 Hz. Participants were instructed to report any observed change in fixation dot color with a button box press. Color changes occurred around once every three seconds. The contrast patterns for the mapping stimuli were first used by Benson et al (2018).

MRI acquisition

Imaging was conducted at the Center for Brain Imaging at New York University using a 3T Siemens Prisma MRI system and a Siemens 64-channel head/neck coil. We acquired functional images with a T2*-weighted multiband echo planar imaging (EPI) sequence with whole-brain coverage (repetition time = 1 s, echo time = 37 ms, flip angle = 68º, 66 slices, 2 x 2 x 2 mm voxels, multiband acceleration factor = 6, phase-encoding = posterior-anterior). We collected

spin echo images with anterior-posterior and posterior-anterior phase-encoding to estimate, and correct for, the susceptibility-induced distortion in the functional EPIs. We also acquired one to three whole-brain T1-weighted MPRAGE 3D anatomical volumes (.8 x .8 x .8 mm voxels) for each of the eight subjects.

MRI processing

All original MRI data (DICOM files) were defaced to anonymize them using *pydeface* (https://github.com/poldracklab/pydeface). The DICOM data were then converted to NIFTI and organized into the Brain Imaging Data Structure format (Gorgolewski et al., 2016) using Heuristic Dicom Converter (Halchenko et al., 2018). The data were then preprocessed using fMRIPrep 20.2.7 (Esteban et al., 2018, 2019), which is based on Nipype 1.7.0 (Gorgolewski et al., 2011, 2018).

Anatomical data preprocessing.

The following sections on anatomical and functional data preprocessing are provided by the fMRIPrep boilerplate text generated by the preprocessed scan output.

Each of the one to three T1w images was corrected for intensity non-uniformity with N4BiasFieldCorrection (Tustison et al., 2010), distributed with ANTs 2.3.3 (Avants et al., 2008). The T1w-reference was then skull-stripped with a Nipype implementation of the antsBrainExtraction.sh workflow (from ANTs), using OASIS30ANTs as target template. Brain tissue segmentation of cerebrospinal fluid, white-matter and gray-matter was performed on the brain-extracted T1w using fast (FSL 5.0.9, (Zhang et al., 2001)). A T1w-reference map was computed after registration of the T1w images (after intensity non-uniformity-correction) using mri_robust_template (FreeSurfer 6.0.1, (Reuter et al., 2010)). Brain surfaces were reconstructed using recon-all (FreeSurfer 6.0.1, (Dale et al., 1999)), and the brain mask estimated previously was refined with a custom variation of the method to reconcile ANTs-derived and FreeSurfer-derived segmentations of the cortical gray-matter of Mindboggle (Klein, 2017).

Functional data preprocessing.

For each of the 12 BOLD runs found per subject (across all tasks and sessions), the following preprocessing was performed. First, a reference volume and its skull-stripped version were generated by aligning and averaging a single-band reference. A B0-nonuniformity map (or fieldmap) was estimated based on two EPI references with opposing phase-encoding directions, with 3dQwarp (Cox and Hyde, 1997). Based on the estimated susceptibility distortion, a

corrected EPI reference was calculated for a more accurate co-registration with the anatomical reference. The BOLD reference was then co-registered to the T1w reference using bbregister (FreeSurfer) which implements boundary-based registration (Greve and Fischl, 2009). Co-registration was configured with six degrees of freedom. Head-motion parameters with respect to the BOLD reference (transformation matrices, and six corresponding rotation and translation parameters) are estimated before any spatiotemporal filtering using mcflirt (FSL 5.0.9, (Jenkinson et al., 2002)). BOLD runs were slice-time corrected to 0.445s (0.5 of slice acquisition range 0s-0.89s) using 3dTshift from AFNI 20160207 (Cox and Hyde, 1997). First, a reference volume and its skull-stripped version were generated using a custom methodology of fMRIPrep. The BOLD time-series were resampled onto the fsnative surface. The BOLD time-series (including slice-timing correction) were resampled onto their original, native space by applying a single, composite transform to correct for head-motion and susceptibility distortions. These resampled BOLD time-series will be referred to as preprocessed BOLD. All resamplings can be performed with a single interpolation step by composing all the pertinent transformations (i.e. head-motion transform matrices, susceptibility distortion correction, and co-registrations to anatomical and output spaces). Gridded (volumetric) resamplings were performed using antsApplyTransforms (ANTs), configured with Lanczos interpolation to minimize the smoothing effects of other kernels (Lanczos, 1964). Non-gridded (surface) resamplings were performed using mri_vol2surf (FreeSurfer).

Many internal operations of fMRIPrep use Nilearn 0.6.2 (Abraham et al., 2014), mostly within the functional processing workflow. For more details of the pipeline, see the section corresponding to workflows in fMRIPrep's documentation.

GLM analyses. From each subject's surface based time series, we used GLMSingle (Prince et al., 2022) to estimate the neural pattern of activity evoked during the 11.5-s delay periods of the main experiment for each trial. GLMSingle is a three step process where 1) an optimal hemodynamic response function (HRF) is fit to each vertex's time series from a bank of 20 HRF functions obtained from the Natural Scenes Dataset via an iterative linear fitting procedure. 2) Noise regressors are computed from the data by identifying noisy vertices defined by negative R^2, deriving noise regressors from this noise pool using principal component analysis, then iteratively removing each noise regressor from all vertices' time series. The optimal number of regressors is determined via a cross-validated R^2 improvement for the task-model. 3) GLMSingle

implements fractional ridge regression as a way to improve robustness of single-trial beta estimates, particularly useful here as our design yields a limited number of trials per target position within each condition.

We constructed our design matrices to have 48 regressors of interest (16 polar angle bins x 3 conditions), with the events modeled as boxcars corresponding to the 11.5 s delay periods. We estimated one model for each participant with these designs in GLMSingle, resulting in single trial estimates for each trial in each surface vertex.

Fitting pRF models

Using the data from the retinotopy session, we fit population receptive field (pRF) models for each vertex on the cortical surface, as described by Himmelberg et al. (2021; section 2.6). In brief, for each surface vertex we fit a circular 2D-Gaussian linear population receptive field (pRF) to the BOLD time series, averaged across identical runs of the bar stimulus. The software was implemented in *Vistasoft* as described in Dumoulin & Wandell (2008), with a wrapper function to handle surface data (https://github.com/WinawerLab/prfVista). The models are parameterized by the Gaussian center (x, y) and standard deviation (σ).

Visual field map definitions

Visual field maps were defined by drawing boundaries at polar angle reversals on each subject's cortical surface using an early version of the visualization tool, cortex-annotate (https://github.com/noahbenson/cortex-annotate), which is built on *neuropythy* software (https://github.com/noahbenson/neuropythy, (Benson and Winawer, 2018). We followed common heuristics to define seven maps spanning early to mid-level visual cortex: V1, V2, V3 (Himmelberg et al., 2021; Benson et al., 2022); hV4 (Winawer and Witthoft, 2015); V3A and V3B (grouped into one ROI, V3ab) and IPS0 (Mackey et al., 2017); and LO1 (Larsson and Heeger, 2006).

We defined experiment-specific regions of interest for each visual field map composed of vertices whose pRF centers were near the target eccentricity and whose variance explained by the pRF model was above 10%. Specifically, we included only those vertices whose pRF centers were within one σ of 7° (the target eccentricity in the experiments). For example, a vertex with pRF center at 6 deg and pRF size (σ) of 1.5 deg would be included, but a vertex with pRF center at 6 deg and pRF size of 0.5 deg would not be included. We imposed the eccentricity restriction

for the purpose of examining polar angle activation profiles, described in the next section. These measures are based on the retinotopy scans only and are therefore independent of the main experiment.

Analyses quantifying perception, working memory, and long-term memory activity

To examine the evoked BOLD response within the 11.5 second delay, we constructed polar-angle activation profiles for each visual field map and each condition (perception, working memory, long-term memory). Some analyses averaged over the delay period. For these analyses, we obtained the response amplitudes from GLMsingle (one beta weight per trial for each surface vertex). The visual field coordinates for each vertex came from pRF mapping. We binned the response amplitudes by the polar angle distance between each vertex's pRF and the target location on that trial. Binning by polar angle distance from the target enabled us to average across trials with different target locations, resulting in an activation profile as a function of distance from the target. This results in a distinct polar angle activation profile for each subject, condition, and visual field map. Prior to averaging across subjects, we normalized each activation profile by dividing by its vector length. To preserve meaningful units, we then rescaled the activation profile by the average of vector lengths across subjects. Visual inspection of the average profiles showed a peak near 0, and negative responses far from 0. We fit the functions with a difference of two von Mises distributions, constrained so that the centers were the same.

A separate procedure was used to derive 2D activation profiles, which include time as well as polar angle. To derive these, we extracted the preprocessed BOLD time series for each vertex on each trial, rather than a single beta weight per trial, expressed as percent signal change from the mean of each scan. Time was sampled at second from 0 s to 14 s in each trial, relative to the start of the delay period. We then computed polar angle activation profiles independently for each time point, using the procedure described above: binning by polar angle distance from the target, averaging across trials within a condition and visual field map, normalizing by the vector length, averaging across subjects, and then re-scaling by the average vector length. This results in 2D activation profiles, which span polar angle and time within a trial.

For temporal analyses these 1D polar angle activation profiles were computed at each timepoint (TR = 1 sec) from the onset of the paired cue to 14 seconds after, resulting in a 2D heatmap of the polar angle activation profiles over time. As in the static activation profiles, the vertices were binned by polar angle distance from the target, and averaged across trials within

each visual map, separately for each condition and subject. This results in a matrix that is polar angle by time. We fit the polar angle profile independently for each time point, again using a difference of Von Mises' distributions, constrained so that the centers were the same.

For both analyses, we estimated the amplitude (trough to peak), the peak location, and full-width-at-half-maximum (FWHM) from each of the difference of Von Mises fits. We bootstrapped across subjects (with replacement) 500 times to obtain 68% and 95% confidence intervals for these location, amplitude, and FWHM parameter estimates.

Temporal analyses

We characterized the time course for each 2D polar angle activation profile by fitting logistic functions to the amplitude estimate at each time point. The logistic fits resulted in estimates of four parameters for each 2D polar angle activation profile: t_0, the rise-time to reach the function's midpoint; L, the upper asymptote of the amplitude estimates; k, the logistic growth rate of the function (i.e. the steepness of the curve); and c a baseline. All four parameters were constrained to have a lower bound of 0; the upper bound was unconstrained for L and c, 15 seconds for t_0, and 5% signal change / second for k. This fit was sufficient for the perception and long-term memory conditions, which generally showed a rise and then a steady response. For working memory, the response rose transiently, and then declined to a lower value. This pattern was accurately captured by fitting a multiplication of two logistic functions rather than a single logistic function. The two functions were constrained to have the same L parameter, but could differ in t_0, c, and k. The parameter bounds for the second logistic function were the same as the first, except where k was constrained to be negative instead of positive. We repeated these logistic fits for each of the bootstraps, computing 68% confidence intervals of the estimated logistic time series.

Saccade analyses

We used 2 EyeLink eye trackers (SR Research, Ottawa, ON, Canada) with a 1000-Hz sampling rate, one in the scanner and one outside the scanner. In the scanner the EyeLink 1000 plus was mounted onto a rig in the magnet bore that holds the projection screen. In the psychophysics room for long-term memory training, we used an EyeLink Tower mount with a 25-mm lens mounted on the EyeLink camera to allow close viewing.

During both the training and the fMRI experiment, saccades were labeled by the EyeLink's default saccade detector, which classifies saccades as eye movements with velocity

and acceleration exceeding 30 deg/s and 8000 deg/s^2, respectively. Saccade responses were collected during the response window at the end of each trial. For all behavioral analyses, the saccade responses used are those which landed nearest the target eccentricity during the saccade response window.

Subjects often make multiple saccades to get to the target. We defined the endpoint as the saccade whose eccentricity was closest to the target eccentricity (7º), irrespective of the polar angle. We then measured the angular distance between this point and the target (ignoring eccentricity). We excluded saccade responses whose eccentricity was less than 3.5° or greater than 12° visual angle from fixation. One subject was removed from the saccade analysis due to technical error with the eye tracker during the scan sessions. Data from three scans from another subject were also excluded due to calibration error.

For comparison between BOLD data and saccades, we divided the saccade data for each subject and each condition into tertiles for counterclockwise, center, and clockwise. We repeated the ROI-level analyses on this split saccade data to obtain 1D polar angle activation profiles for each ROI, condition, subject, and tertile.

Resampling Statistics

We used the bootstrapped data from each analyses to make inferences on spatial tuning properties as a function of condition and of other trial-level factors. Statistics reported are computed using the bootstrapped data. To assess our main claims, we report the mean and confidence intervals from the bootstrapped data. For comparisons between a measurement and a fixed value, we report the 95% CI. For comparisons between two estimates, we report the 68% CIs.

To assess the relationship between a map's position in the visual hierarchy and spatial tuning width, we assigned each map an ordinal value according to its relative position: 1, V1; 2, V2; 3, V3; 4, hV4, LO1, V3A/B; 5, IPS0. We fit a line of tuning width vs ordinal position for each bootstrap, generating a distribution of slope means for each condition. The rest of our analyses took the form of computing the differences of mean effects between conditions, both for individual visual maps and for early (V1-V3) vs. later (hV4-IPS0) visual cortex.

Software

Data visualization, model fitting, and statistical quantification for all analyses described in this paper were made using matplotlib 3.5.2 (Hunter, 2007), nibabel 3.2.2 (Brett et al., 2022), pandas 1.4.2 (The pandas development team, 2024), scikit-learn 1.0.2 (Pedregosa et al., 2011), scipy 1.8.0 (Virtanen et al., 2020), and seaborn 0.11.2 (Waskom, 2021).

Results

We tested how the spatial tuning of cortical visual representations is shaped by viewing a peripheral target, maintaining it in working memory, or retrieving it from long-term memory. We also tested how the cortical representations during memory relate to memory-guided saccades.

Working memory and long-term memory evoke spatially tuned responses in visual cortex

We first ask whether activation in early visual cortex during memory is spatially tuned. To assess this, we parameterized the sensory representations generated during the delay period by remapping GLM estimates of brain activity from the cortical surface to visual space (Figure 3, left). We then computed polar angle activation profiles to capture the spatial tuning of sensory representations generated during the delay period (Figure 3, right), and estimated their amplitude, peak location, and tuning width. We compared these estimates between perception, working memory, and long-term memory conditions.

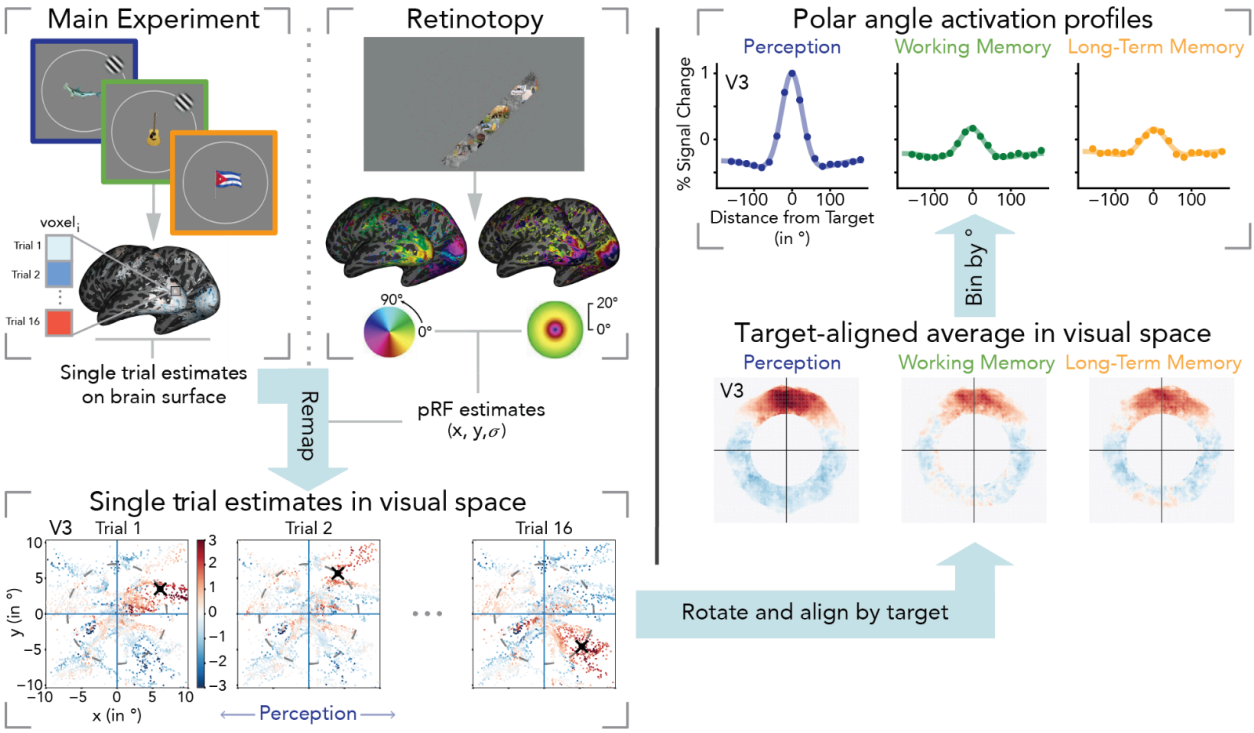

Figure 3. Target-aligned averages of sensory representations in visual space.

Single trial beta estimates for each vertex on the brain surface are obtained from the main task's delay period. In addition, a separate retinotopic mapping procedure is used to obtain population receptive field (pRF) estimates for each vertex. We use the pRF estimates to remap the single trial estimates on the brain surface to single trial estimates in visual space, that are then rotated and aligned by trial target position. The pRF estimates are also used to define retinotopic maps for seven regions of interest across visual cortex, and to restrict voxels to those with pRF centers near the target eccentricity. Normalizing and averaging the aligned trial estimates for each map yields target-aligned averages in visual space for each condition. A 2D visualization of a subject's target-aligned average is shown above for visual map V3. Beta estimates are binned by polar angle distance from target and fit to a difference of Von Mises functions to produce polar angle activation profiles. This mapping of evoked BOLD response as a function of polar angle distance from the stimulus captures the spatial tuning profile of cortical responses during perception, working memory, and long-term memory. Panels in this figure and all following figures can be reproduced using code contained in the /paper/figures folder at https://github.com/rfw256/Woodry_2024_Cortical-tuning-of-visual-memory/. The panels 'Single trial estimates' and 'target-aligned averages' are generated using *fig3_03-04-2024.py*.

The clearest difference between perception and the two memory conditions is that the BOLD amplitude is much larger during perception. This difference is particularly evident in earlier visual maps V1-V4. While the amplitudes were lower during memory, they were all positive, ranging from 0.25% to 0.5% percent signal change across visual maps (Figure 4B. middle). The 95% confidence interval did not overlap 0% BOLD response in any ROI.

The memory activation profiles, like the perception activation profiles, were spatially tuned to the stimulus. Specifically, for all 3 conditions and all 7 visual field maps, peak location estimates were centered around 0° relative to the stimulus angle (Figure 4B, left). Polar angle activation profiles which peak at 0° indicate accurate tuning to the true target location. Tuning to the target location is of course expected in the perception condition. But remarkably, we even see tuning to the target location in the earliest visual map, V1, during both memory conditions. (The confidence intervals from all three conditions include 0º.) Tuning in the two memory conditions confirms prior work showing engagement of visual areas, including primary visual cortex, during memory (Breedlove et al., 2020; Favila et al., 2022; Vo et al., 2022).

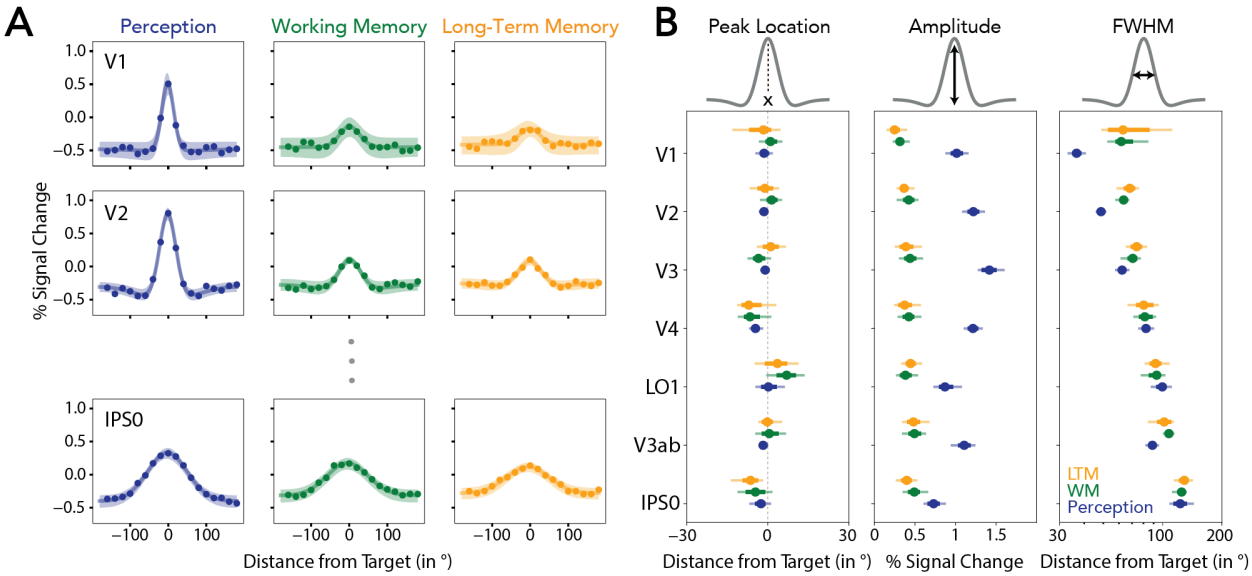

Figure 4. Memory has broader spatial tuning than perception in earlier visual cortex.

A) Polar angle activation profiles fit to brain activity during 11.5 second presentation/delay period. Highlighted regions indicate 68% confidence intervals bootstrapped across subjects. Perception working memory, and long-term memory brain activity show spatial tuning to target locations across visual ROIs. B) Spatial tuning metrics obtained from polar angle activation profiles across visual ROIs. The dots are means across bootstraps. The thick shading is the 68% confidence interval. The thin shading is the 95% confidence interval. Dotted line in the leftmost panel represents the target location (0°). The panels here are generated using *fig4_03-04-2024.py*.

Memory shows broader spatial tuning than perception in early visual cortex

The lower amplitude but similar peak location during memory compared to perception suggests that memory responses might be the same as perception except for a scale factor. This turns out to not be correct. We find that instead of memory responses looking like perception responses

(up to a scale factor), the memory responses show a difference in tuning width, which cannot be achieved by simply increasing or decreasing the response amplitude.

During perception trials, tuning widths increased sharply from early to later visual maps. To quantify this tendency, we assigned each map an ordinal value based on its estimated position in the visual hierarchy: 1, V1; 2, V2; 3, V3; 4, hV4, LO1, V3A/B; 5, IPS0. We fit a line of tuning width vs ordinal position, and find that for perception, the tuning width increases about 21º per position in the hierarchy (slope = 20.9°, CI [17.9, 23.9]), consistent with what is known about receptive fields increasing in size from early to later visual areas (Smith et al., 2001); (Dumoulin and Wandell, 2008). In contrast, there was little increase in tuning width across visual maps during both forms of memory, with a slope nearly half that measured during perception (working memory: slope = 14°, CI [10.1, 16.9]; long-term memory: slope = 13.2°, CI [4.1, 18.7]). Because the tuning widths were so similar for long-term memory and working memory, we compared the average of the two memory conditions to perception. This steeper slope for perception than memory was robust (diff = 7.3º, CI [3.9, 13.5]).

In early visual maps V1-V3, spatial tuning during memory was broader than during perception (diff = 19°, CI [13.1, 28]), especially in V1 (diff = 30.7°, CI [14.2, 59.3]; Figure 4B, right). In later maps – hV4, LO1, V3AB, and IPS0 – tuning widths were about the same in perception, working memory, and long-term memory (diff = 2.8°, CI [-3.6, 8.3]). Therefore, the memory responses are not just a scaled down version of the perceptual responses, but rather show broader tuning in earlier maps.

Because the only differences between conditions are how visual information is accessed, this confirms our hypothesis that differences in how information enters visual cortex shapes the spatial tuning of the subsequent representations.

Distinct spatial tuning profiles emerge over time

The analysis above showed that the routing of stimulus information to visual cortex –feedforward in perception vs top-down in memory– affects the spatial representation. Here we ask how the routing affects the temporal dynamics of the response. Our expectation is that long-term memory responses will be slowest because the responses must be generated entirely internally (no stimulus is viewed). We can also ask whether working memory and long-term memory signals show the same tendency to sustain over the delay. To compare the temporal dynamics across conditions, we averaged the BOLD time series throughout the delay period

across trials, binned by polar angle distance from the target (Figure 5a). We then computed the peak amplitude at each time point, and fit a rising logistic function (perception and long-term memory) or the product of a rising and falling logistic (working memory) across the time series (Figure 5b). The logistic product captures both the transient response evoked by the target at the beginning of the working memory delay period and the decay to a lower, sustained, activation for the remainder of the delay period. Because we fit the working memory response using a different function, we do not compare logistic parameters to the other two conditions. We make three observations that distinguish perception from long-term memory from working memory.

First, the working memory responses show a clear transition from a stimulus-driven transient, peaking at about 4 to 5 seconds after the cue, to a lower sustained signal. This is expected because the target stimulus is briefly shown prior to the delay during working memory.

Second, we find a general tendency for slower rise-time in long-term memory than perception across visual maps, with longer rise times in 6 of 7 maps, more prominent in V2 and V4 (Figure 5c). This is consistent with memory responses arising later due to the sluggishness of feedback. Moreover, the rise times are much more variable in memory, as expected from responses that are internally and effortfully generated (memory) rather than from external, stimulus-triggered responses (perception). Because of this variability, the rise-time is not estimated precisely in long-term memory. Hence more data, either from more subjects or more trials, would be needed to quantitatively test the claim.

Last, there is a slight tendency for the long-term memory and the working memory responses to cross over toward the end of the trial, as working memory signals appear to decline and long-term signals do not (Figure 5d). This pattern could arise from continued retrieval of the stimulus throughout the long-term memory trials, and decreasing strength of the memory trace during the working memory delay. Alternatively, the decrease during the working memory delays could be a post-stimulus undershoot in the hemodynamic response. Experiments with longer and variable delays are needed to disentangle these possibilities.

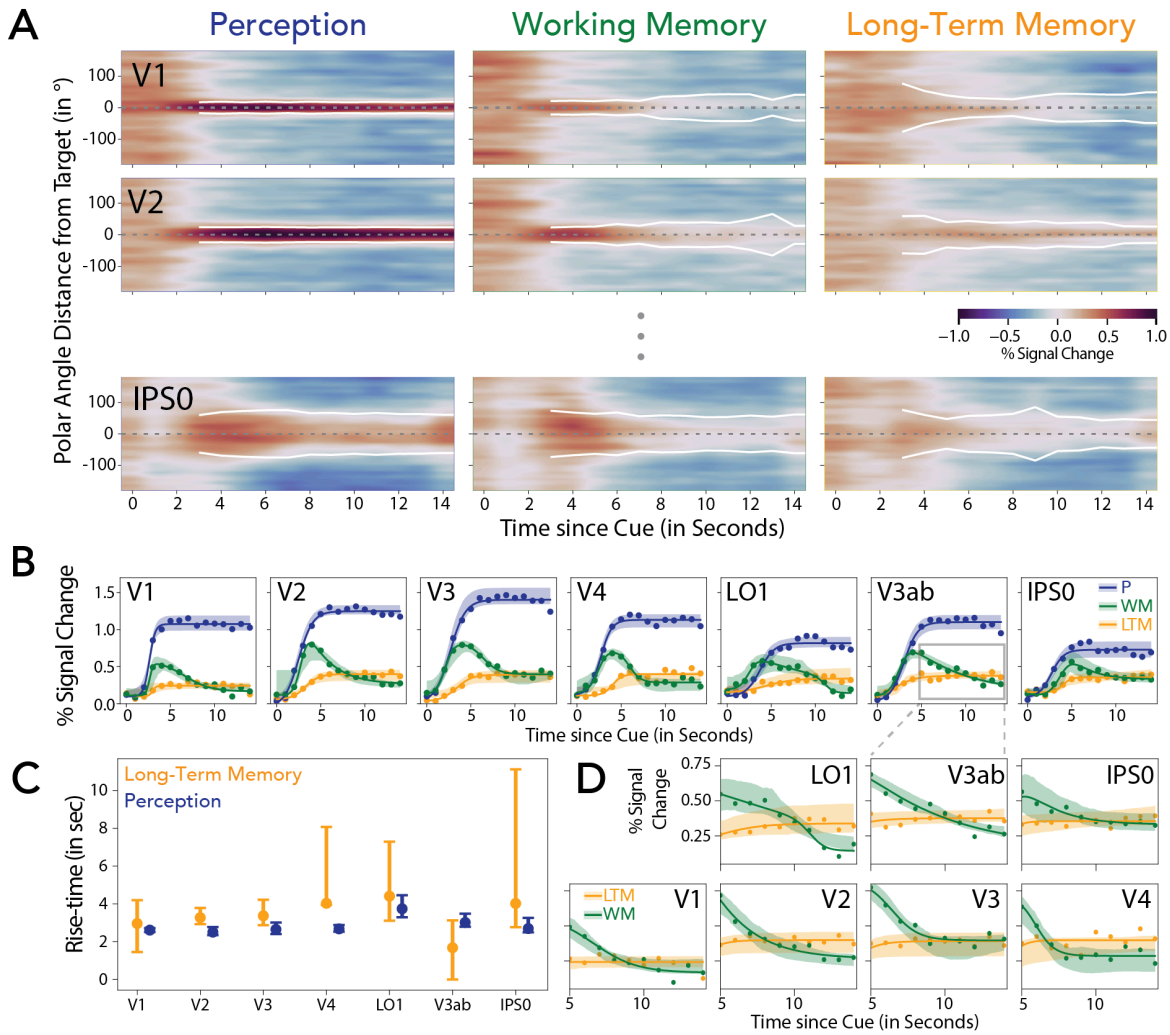

Figure 5. Perception, working memory, and long-term memory have distinct tuning profiles over time.

A) BOLD time course during the delays periods across visual cortex. Target location is plotted as a dashed gray line at 0°. Full-width at half-maximum estimates are plotted as white lines above and below the target location line, starting from 3 seconds after the onset of the delay. Memory-based sensory representations vary in their tuning over time, whereas perceptually evoked responses remain stable. B) Logistic functions fitted to amplitude estimates over time. Highlighted regions represent the 68% confidence interval of the logistic fits, bootstrapped across subjects. C) Comparison of rise-times to the inflection point of logistic fits between perception (blue) and long-term memory (orange) across visual maps. Error bars represent 68% confidence intervals bootstrapped across subjects. Long-term memory responses take longer to rise than perceptually evoked responses. D) Zoomed in version of panel B showing the comparison of response amplitudes near the end of the delay period between working memory (green) and long-term memory (orange) across visual maps. Error bars represent 68% confidence intervals bootstrapped across subjects. There is a crossover in working/long-term memory amplitude estimates in the latter part of the delay period for most visual maps. The panels here are generated using *fig5_03-04-2024.py*.

Errors in cortical tuning aligned with errors in memory-guided behavior

The results above demonstrate that visual cortex is engaged in a retinotopically specific way during working and long-term memory. These results do not, however, indicate whether these retinotopic representations are relevant for behavior. This is an important and open question in the field, with some reports claiming the representations are linked to behavior (Bone et al., 2019, 2020; Hallenbeck et al., 2021; Li et al., 2021) and some questioning their relevance (Xu, 2017). If the visual cortex representations are relevant for behavior, we expect alignment between the memory-driven cortical responses during the delay and subsequent saccade responses. Here, we took advantage of trial variability to test whether the peak location of cortical responses during the delay aligned with the direction of saccade error.

To test the alignment between cortical and saccade responses, we split the trials into three groups by their saccade error (Figure 6b), those trials with saccades near the targets, counterclockwise of the targets, or clockwise of the targets, with angular thresholds set to make the three bins contain equals number of trials. We then repeated our spatial tuning analyses separately for the clockwise and counterclockwise trials to compute the estimates of peak location during the delay period (Figure 6c). To reduce the number of comparisons, we defined a new region of interest as the union of all 7 maps. For this large ROI, there was a strong link between neural tuning and saccade error in long-term memory: the peak estimates for trials with clockwise saccades were 19.5º more clockwise than the trials with counterclockwise saccades (diff = -19.5°, CI = [-34.2, -9.2]). In each of the separate maps, the same general pattern is found in long-term memory. It is particularly pronounced in V3 (diff = -20.1°, CI [-29.9, -8.9]), V4 (diff = -25.9°, CI [-65.8, -1.5]), V3ab (diff = -14.7°, CI [-36.5, 2.7]), and IPS0 (diff = -39.4°, CI [-62.4, 1.7]). In no map does the tuning go in the opposite direction of the saccades.

In contrast, for working memory, there was no systematic relationship between the peak estimates of spatial tuning and saccade direction (diff = 1.7°, CI [-6.9, 11.6]). And there was at best a small effect in perception (diff = -2.5°, CI [-6.0, 0.7]). This could be due to a restricted range of saccade errors in perception, and to the possibility that when there are errors in working memory, they develop gradually over the trial, as the representation drifts away from the viewed target. Since the analysis of peak location pools across the whole delay, this would diminish the impact of biases near the ends of the trials.

The long-term memory results suggest that information reconstructed in visual maps is likely used for memory-guided behavior.

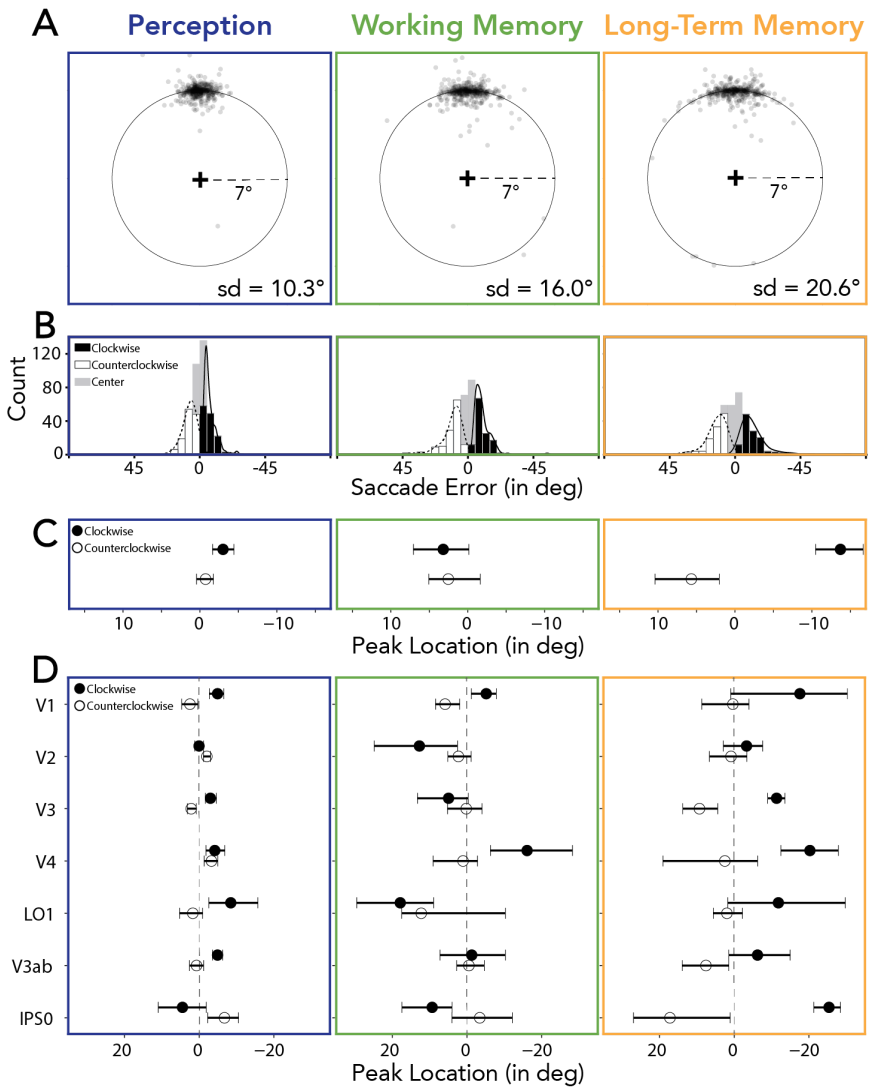

Figure 6. Saccade errors align with cortical tuning in long-term memory.

A) Saccade response spread in visual space. We aligned saccade responses to the target location by subtracting their polar angle distance from target, and added 90° to align to the vertical axis. Saccade responses during memory had greater error along polar angle than during perception, consistent with the broader spatial tuning during memory within early visual maps. Standard deviation of saccade errors in polar angle are shown in the bottom right. B) Histograms of saccade error along polar angle. We grouped trials into tertiles: clockwise (black), center (gray), or counterclockwise (white) to target polar angle location. The tertiles were defined for each subject and conditions. C) Peak location estimates for clockwise/counterclockwise trial groups. We fit polar angle activation profiles to the clockwise (black circles) and counterclockwise (white circles) groups. Error bars represent the bootstrapped 68% confidence intervals. D) Peak location estimates for clockwise/counterclockwise groups for each measured visual map. Error bars represent 68% confidence intervals. The panels here are generated using *fig6_04-08-2024.py*.

Discussion

We measured the spatial tuning of cortical responses during perception, working memory, and long-term memory. We find that both long-term and working memory scale the spatial tuning of responses in early but not late visual field maps and that the temporal dynamics during the delay differ across the three conditions. We further demonstrate the importance of memory-driven visual cortex activations by showing that they correlate with behavior: trial-to-trial variation in the spatial tuning of cortical responses during long-term memory is related to the accuracy of the subsequent oculomotor responses.

Spatial tuning in visual cortex during long-term memory differs from perception

Long-term memory retrieval of information from hippocampus is reported to reinstate patterns of cortical activity evoked during encoding (Wheeler et al., 2000; Rugg et al., 2008; Johnson et al., 2009; Pearson et al., 2015). This 'cortical reinstatement' is supported by neurally-inspired models of episodic memory (Damasio, 1989; McClelland et al., 1995; Rolls, 2000; Teyler and Rudy, 2007). These theories do not imply that memory representations are identical to perceptual representations. For example, it is widely found that memory responses are noisier than perceptual responses. One possibility is that neural responses during memory reinstatement are like those during perception, except for decreased signal-to-noise level (Wheeler et al., 2000; Rugg et al., 2008; Johnson et al., 2009; Pearson et al., 2015).

We find partial support for such reinstatement. On the one hand, spatial tuning evoked by long-term memory retrieval peaks at the expected locations in multiple visual field maps, as early as V1. This supports the claim that the reinstatement of sensory information during memory is stimulus specific. On the other hand, the long-term memory responses in earlier visual maps (V1-V3) were more broadly tuned than perceptual responses, consistent with two prior reports (Breedlove et al., 2020; Favila et al., 2022). The systematic difference in tuning width between memory and perception demonstrates that memory driven-activity is not simply a reinstatement of perceptual responses. A likely explanation lies in the architecture of visual cortex: during feedback, earlier maps inherit the lower spatial precision characteristic of later visual maps where neurons have large receptive fields. Lower precision in later stages is expected from a hierarchical encoding process with increasing levels of tolerance to size and position of stimuli (Ito et al., 1995; Kay et al., 2013): the greater size and position tolerance in

later visual maps reflects a loss of position information likely not recoverable during retrieval. This information loss puts a limit on the precision of cortical reinstatement, independent of the fact memory responses also tend to have weaker (i.e., lower amplitude) signals.

How much information is lost in long-term memory?

In addition to architectural constraints on the precision of top-down generated signals, the precision may also depend on task demands. A comparison between our results and those of Favila et al. supports this idea. We replicated Favila et al. (2022) in that both studies showed broader tuning for memory than perception in early visual maps. However, the decreased precision during memory was more pronounced in Favila's study than in this one (Figure 7a). This difference is not likely due to measurement noise, difference in eccentricity (4º vs 7º), or differences in analysis. We can infer this because the tuning width estimates from both studies are the same in all maps during perception, and in later maps during memory - the only differences between them being long-term memory in V1 to V3. We speculate that the difference arises from encoding demands. The studies differed in the number of remembered targets and their spacing: every 90 deg in Favila et al (four learned targets), compared to an average of 22.5 deg here (16 targets). The narrow spacing in our study likely required more precise encoding, leading to narrow tuning in V1-V3 during memory. A post-hoc analysis supports this. Due to the random placement of stimulus locations within polar angle bins, the spacing between our stimuli varied (Figure 7c). For some stimuli, the spacing was just a few degrees. For others it was over 30 degrees. Although two stimuli were never simultaneously present, the training protocol required saccade responses to be closer to the correct target than to any other target. Hence nearby targets required greater precision. In V1-V3, tuning was narrower for targets with near neighbors than targets with far neighbors (Figure 7b). The tuning width for far-spaced targets is more similar to Favila et al's results. The implication is that increased competition during memory training invokes a greater demand for spatial precision, resulting in narrower tuning during recall. Further studies are needed to test whether this effect is driven by hippocampal pattern separation (Bakker et al., 2008; Yassa and Stark, 2011).

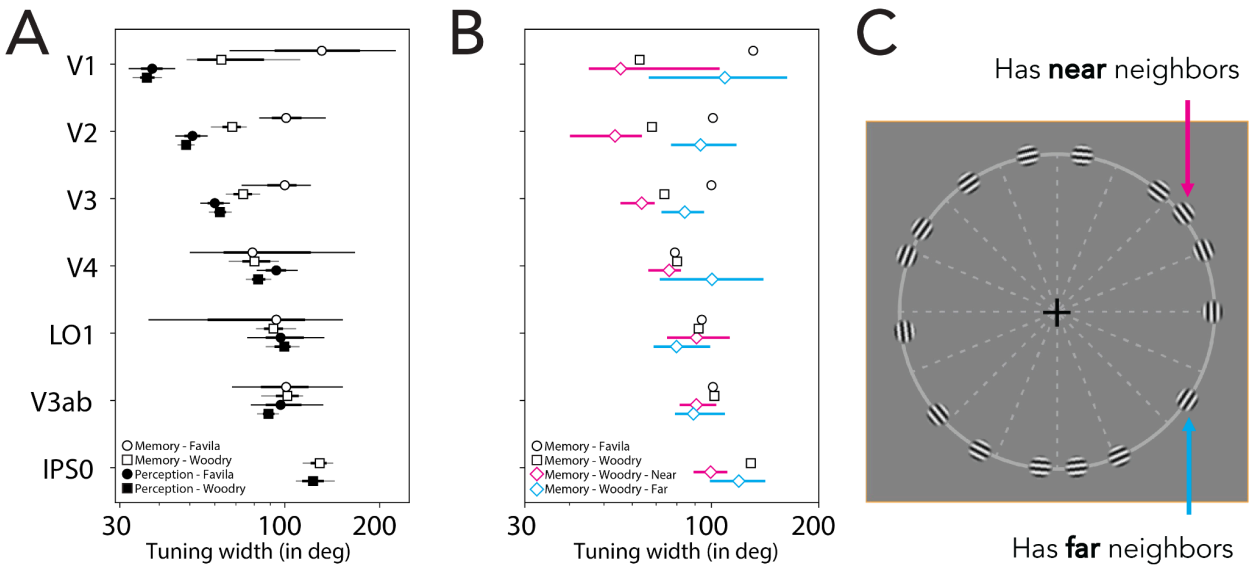

Figure 7. Systematic differences in spatial tuning of cortical responses.

A) Comparison of spatial tuning (FWHM) estimates across measured visual maps between Favila et al. (2022; circles) and our study (squares). Both studies show large agreement in their spatial tuning estimates during the perception condition (filled), and in later visual maps V4-IPS0 during long-term memory retrieval (unfilled). During long-term memory, tuning widths differ in early visual maps V1-V3. The main differences between ours and the Favila study's memory conditions are in the spacing of the stimuli; their study required the retrieval of four targets equally spaced around the visual field, whereas ours involved the retrieval of 16 targets that varied in their separation. B) Spatial tuning width estimates across visual maps, split by whether targets were spaced near to (pink diamonds; spacing: 1.7° - 12°, mean: 7.7°, std: 3.3°) or spaced far from (blue diamonds; spacing: 19.2° - 31.8°, mean: 23°, std: 3.4°) each other. Also plotted are the mean tuning width estimates from both the Favila et al. study (circles) and this study (squares). The tuning widths from far-spaced targets in our study more closely resemble the broader tuning observed in the Favila et al. study, suggesting the differences between mean width estimates between our studies is in large part due to the increased precision required of more closely spaced targets. C) Example of this study's long-term memory retrieval targets. 16 Target stimuli, each sampled from within 22.5° bins, vary in their separation from other targets. Some targets were spaced far apart, while others were spaced near each other, meaning those spaced near to each other likely required more precision during later recall. This is reflected in the narrower spatial tuning observed during these trials than for the "far-spaced" trials. Panels A and B are generated using *fig7_03-04-2024.py*.

Spatial tuning in visual cortex during working memory is similar to long-term memory

Excluding the early part of the delay period, the responses during working memory were much more similar to long-term memory than perception in both amplitude and tuning width. This may be due to a partially overlapping mechanism for both types of memory: As information is fed back from top-down sources, spatial precision is limited by relatively coarse tuning in the mid- to

high-level areas through which the information is routed. What limits the precision appears to be not whether the stimulus was recently viewed (working memory) or retrieved from long-term storage, but rather that the stimulus representation is maintained in the absence of perceptual input. In contrast, in psychophysical experiments where a stimulus is viewed dozens to hundreds of times, highly precise low-level (eye- and orientation-specific) sensory information can be maintained for up to several minutes (Ishai and Sagi, 1995, 1997).

Surprisingly, responses in V1-V3 were more precise in some long-term memory trials (near spaced stimuli, see fig. 7b) than working memory, despite the longer delay between encoding and retrieval for long-term memory. This is likely because during our long-term memory experiment, subjects viewed the stimuli many times (about 60 times each during pre-scan training), and for near-spaced stimuli, needed to make fine distinctions. Recent work shows that when stimuli are highly learned in long-term memory, the working memory of these stimuli becomes more precise (Miller et al., 2022), similar to the relatively precise representations in V1-V3 for our near spaced long-term memory stimuli.

The temporal response profile in visual cortex during working memory differs from long-term memory

While the average spatial tuning across the delay was similar between working and long-term memory, the temporal profiles were distinct. The most obvious difference is the large initial transient seen in working memory due to the presence of the stimulus at the onset of the delay. But even after this transient, we observed a greater amplitude decline in working memory than in long-term memory for many visual maps. One possibility is that the working memory signal degrades over time without continuous input via the eyes or retrieval from long-term memory. There is some support for time-dependent decay in working memory (Brown, 1958; Cowan et al., 1997; Ricker and Cowan, 2010; Mercer and Barker, 2020), although when, whether, and how this happens remains controversial (Lewandowsky et al., 2004; Zhang and Luck, 2009; Ricker et al., 2016). An alternative, garnering increasing recent support, is that the stimulus-driven response is instead transformed into a more abstract representation throughout the delay period (Kwak and Curtis, 2022; Li and Curtis, 2023). Such transformation will yield weaker amplitude in analyses such as ours that do not explicitly model and measure the abstract representation.

We observed no decline in the response during the long-term memory delay: once the responses peaked, the amplitude remained steady, similar to during perception but unlike during

working memory. This suggests that during long-term memory, the retrieved perceptual information is continuously routed to early visual maps. This is consistent with the finding that visual working memory representations are less tightly linked to detailed image features compared to representations in long-term memory (Schurgin and Flombaum, 2018).

Decoding vs encoding approaches to studying memory representations

Our use of retinotopic models allowed us to quantify and compare sensory activation during perception, working memory, and long-term memory. We did this by measuring spatial tuning functions in all three conditions. This differs from the decoding approach, often used for stimulus orientation or other stimulus features which do not vary systematically at the mm to cm scale of fMRI measurement. Decoding results of fMRI responses show reduced accuracy in memory or imagery compared to perception (Naselaris et al., 2015). Such results have been coupled to the idea that memory or imagery representations are like weak forms of perception, similar to viewing stimuli at low contrast (Pearson, 2019). But the decrease in decoding accuracy can arise from a decrease in signal-to-noise ratio (expected from a weaker stimulus) or a decrease in precision, or a combination. Our stimulus manipulation allowed us to quantify the precision independent of amplitude, and showed that both are lower in memory, and that reduced precision varies with task demands and map location in the visual hierarchy. Separating precision of the representation from the signal-to-noise ratio is difficult for stimulus features whose representations vary over a sub mm scale, such as orientation, despite recent advances in studying their encoding (Roth et al., 2018; Sprague et al., 2018; Gardner and Liu, 2019). In contrast, it is straightforward in the spatial domain, making spatial manipulations a useful tool for comparing neural representations across conditions. But spatial manipulations are not just a tool. Space is a fundamental part of visual representations (Wandell and Winawer, 2011) and its role in both perception and memory disorders remains a central topic in cognitive neuroscience (Bisiach and Luzzatti, 1978; Farah, 2003).

References

Abraham A, Pedregosa F, Eickenberg M, Gervais P, Mueller A, Kossaifi J, Gramfort A, Thirion B, Varoquaux G (2014) Machine learning for neuroimaging with scikit-learn. Front Neuroinform 8:14.

Albers AM, Kok P, Toni I, Dijkerman HC, de Lange FP (2013) Shared representations for working memory and mental imagery in early visual cortex. Curr Biol 23:1427–1431.

Avants BB, Epstein CL, Grossman M, Gee JC (2008) Symmetric diffeomorphic image registration with cross-correlation: evaluating automated labeling of elderly and neurodegenerative brain. Med Image Anal 12:26–41.

Bakker A, Kirwan CB, Miller M, Stark CEL (2008) Pattern separation in the human hippocampal CA3 and dentate gyrus. Science 319:1640–1642.

Benson NC, Winawer J (2018) Bayesian analysis of retinotopic maps. Elife 7 Available at: http://dx.doi.org/10.7554/eLife.40224.

Benson NC, Yoon JMD, Forenzo D, Engel SA, Kay KN, Winawer J (2022) Variability of the Surface Area of the V1, V2, and V3 Maps in a Large Sample of Human Observers. J Neurosci 42:8629–8646.

Bisiach E, Luzzatti C (1978) Unilateral neglect of representational space. Cortex 14:129–133.

Bone MB, Ahmad F, Buchsbaum BR (2020) Feature-specific neural reactivation during episodic memory. Nat Commun 11:1945.

Bone MB, St-Laurent M, Dang C, McQuiggan DA, Ryan JD, Buchsbaum BR (2019) Eye Movement Reinstatement and Neural Reactivation During Mental Imagery. Cereb Cortex 29:1075–1089.

Bosch SE, Jehee JFM, Fernández G, Doeller CF (2014) Reinstatement of associative memories in early visual cortex is signaled by the hippocampus. J Neurosci 34:7493–7500.

Breedlove JL, St-Yves G, Olman CA, Naselaris T (2020) Generative Feedback Explains Distinct Brain Activity Codes for Seen and Mental Images. Curr Biol 30:2211–2224.e6.

Brodeur MB, Dionne-Dostie E, Montreuil T, Lepage M (2010) The Bank of Standardized Stimuli (BOSS), a new set of 480 normative photos of objects to be used as visual stimuli in cognitive research. PLoS One 5:e10773.

Brown J (1958) Some Tests of the Decay Theory of Immediate Memory. Q J Exp Psychol 10:12–21.

Brett, M., Markiewicz, C. J., Hanke, M., Côté, M.-A., Cipollini, B., McCarthy, P., Jarecka, D., Cheng, C. P., Halchenko, Y. O., Cottaar, M., Larson, E., Ghosh, S., Wassermann, D., Gerhard, S., Lee, G. R., Wang, H.-T., Kastman, E., Kaczmarzyk, J., Guidotti, R., … freec84.

(2022). nipy/nibabel: 3.2.2 (3.2.2). Zenodo. https://doi.org/10.5281/zenodo.6617121

Cowan N, Saults JS, Nugent LD (1997) The role of absolute and relative amounts of time in forgetting within immediate memory: The case of tone-pitch comparisons. Psychon Bull Rev 4:393–397.

Cox RW, Hyde JS (1997) Software tools for analysis and visualization of fMRI data. NMR Biomed 10:171–178.

Curtis CE, D'Esposito M (2003) Persistent activity in the prefrontal cortex during working memory. Trends Cogn Sci 7:415–423.

Curtis CE, Sprague TC (2021) Persistent Activity During Working Memory From Front to Back. Front Neural Circuits 15:696060.

Dale AM, Fischl B, Sereno MI (1999) Cortical surface-based analysis. I. Segmentation and surface reconstruction. Neuroimage 9:179–194.

Damasio AR (1989) Time-locked multiregional retroactivation: a systems-level proposal for the neural substrates of recall and recognition. Cognition 33:25–62.

D'Esposito M, Postle BR (2015) The cognitive neuroscience of working memory. Annu Rev Psychol 66:115–142.

Dumoulin SO, Wandell BA (2008) Population receptive field estimates in human visual cortex. Neuroimage 39:647–660.

Esteban O, Blair R, Markiewicz CJ, Berleant SL, Moodie C, Ma F, Isik AI, Erramuzpe A, Kent JD, Goncalves M, Others (2018) fMRIPrep. Software. Zenodo.

Esteban O, Markiewicz CJ, Blair RW, Moodie CA, Isik AI, Erramuzpe A, Kent JD, Goncalves M, DuPre E, Snyder M, Oya H, Ghosh SS, Wright J, Durnez J, Poldrack RA, Gorgolewski KJ (2019) fMRIPrep: a robust preprocessing pipeline for functional MRI. Nat Methods 16:111–116.

Farah MJ (2003) Disorders of visual-spatial perception and cognition. Clinical neuropsychology 4 Available at: https://books.google.com/books?hl=en&lr=&id=MT_RCwAAQBAJ&oi=fnd&pg=PA152&dq=Farah+and+Epstein+2003&ots=-pQhhikHZo&sig=uSlYDe0ibIeT4wfc7toVl8DFHmM.

Favila SE, Kuhl BA, Winawer J (2022) Perception and memory have distinct spatial tuning properties in human visual cortex. Nat Commun 13:5864.

Gardner JL, Liu T (2019) Inverted Encoding Models Reconstruct an Arbitrary Model Response, Not the Stimulus. eNeuro 6 Available at: http://dx.doi.org/10.1523/ENEURO.0363-18.2019.

Gorgolewski K, Burns CD, Madison C, Clark D, Halchenko YO, Waskom ML, Ghosh SS (2011) Nipype: a flexible, lightweight and extensible neuroimaging data processing framework in

python. Front Neuroinform 5:13.

Gorgolewski KJ et al. (2016) The brain imaging data structure, a format for organizing and describing outputs of neuroimaging experiments. Scientific Data 3:1–9.

Gorgolewski KJ, Nichols T, Kennedy DN, Poline J-B, Poldrack RA (2018) Making replication prestigious. Behav Brain Sci 41:e131.

Greve DN, Fischl B (2009) Accurate and robust brain image alignment using boundary-based registration. Neuroimage 48:63–72.

Halchenko, Y. O. et al. Open Source Software: Heudiconv. Zenodo, https://doi.org/10.5281/zenodo.1306159 (2018).

Hallenbeck GE, Sprague TC, Rahmati M, Sreenivasan KK, Curtis CE (2021) Working memory representations in visual cortex mediate distraction effects. Nat Commun 12:4714.

Harrison SA, Tong F (2009) Decoding reveals the contents of visual working memory in early visual areas. Nature 458:632–635.

Himmelberg MM, Kurzawski JW, Benson NC, Pelli DG, Carrasco M, Winawer J (2021) Cross-dataset reproducibility of human retinotopic maps. Neuroimage 244:118609.

Hunter J, "Matplotlib: A 2D Graphics Environment", Computing in Science & Engineering, vol. 9, no. 3, pp. 90-95, 2007

Inouye T (1909) Die Sehstörungen bei Schussverletzungen der kortikalen Sehsphäre: nach Beobachtungen an Verwundeten der letzten japanischen Kriege. Engelmann.

Ishai A, Sagi D (1995) Common mechanisms of visual imagery and perception. Science 268:1772–1774.

Ishai A, Sagi D (1997) Visual imagery facilitates visual perception: psychophysical evidence. J Cogn Neurosci 9:476–489.

Ito M, Tamura H, Fujita I, Tanaka K (1995) Size and position invariance of neuronal responses in monkey inferotemporal cortex. J Neurophysiol 73:218–226.

Jenkinson M, Bannister P, Brady M, Smith S (2002) Improved optimization for the robust and accurate linear registration and motion correction of brain images. Neuroimage 17:825–841.

Johnson JD, McDuff SGR, Rugg MD, Norman KA (2009) Recollection, familiarity, and cortical reinstatement: a multivoxel pattern analysis. Neuron 63:697–708.

Kay KN, Winawer J, Mezer A, Wandell BA (2013) Compressive spatial summation in human visual cortex. J Neurophysiol 110:481–494.

Klein A (2017) Mindboggle-101 templates (unlabeled images from a population of brains). Harvard Dataverse.

Kosslyn SM, Thompson WL, Kim IJ, Alpert NM (1995) Topographical representations of mental images in primary visual cortex. Nature 378:496–498.

Kwak Y, Curtis CE (2022) Unveiling the abstract format of mnemonic representations. Neuron 110:1822–1828.e5.

Lanczos C (1964) Evaluation of Noisy Data. Journal of the Society for Industrial and Applied Mathematics Series B Numerical Analysis 1:76–85.

Larsson J, Heeger DJ (2006) Two retinotopic visual areas in human lateral occipital cortex. J Neurosci 26:13128–13142.

Lewandowsky S, Duncan M, Brown GDA (2004) Time does not cause forgetting in short-term serial recall. Psychon Bull Rev 11:771–790.

Li H-H, Curtis CE (2023) Neural population dynamics of human working memory. Curr Biol 33:3775–3784.e4.

Li H-H, Sprague TC, Yoo AH, Ma WJ, Curtis CE (2021) Joint representation of working memory and uncertainty in human cortex. Neuron 109:3699–3712.e6.

Mackey WE, Winawer J, Curtis CE (2017) Visual field map clusters in human frontoparietal cortex. Elife 6 Available at: http://dx.doi.org/10.7554/eLife.22974.

McClelland JL, McNaughton BL, O'Reilly RC (1995) Why there are complementary learning systems in the hippocampus and neocortex: insights from the successes and failures of connectionist models of learning and memory. Psychol Rev 102:419–457.

Mercer T, Barker E (2020) Time-dependent forgetting in visual short-term memory. J Cogn Psychol 32:391–408.

Miller JA, Tambini A, Kiyonaga A, D'Esposito M (2022) Long-term learning transforms prefrontal cortex representations during working memory. Neuron 110:3805–3819.e6.

Naselaris T, Olman CA, Stansbury DE, Ugurbil K, Gallant JL (2015) A voxel-wise encoding model for early visual areas decodes mental images of remembered scenes. Neuroimage 105:215–228.

The pandas development team. (2024). pandas-dev/pandas: Pandas (v2.2.1). Zenodo. https://doi.org/10.5281/zenodo.10697587

Pearson J (2019) The human imagination: the cognitive neuroscience of visual mental imagery. Nat Rev Neurosci 20:624–634.

Pearson J, Naselaris T, Holmes EA, Kosslyn SM (2015) Mental Imagery: Functional Mechanisms and Clinical Applications. Trends Cogn Sci 19:590–602.

Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, Blondel M, Louppe G,

Prettenhofer P, Weiss R, Weiss RJ, Vanderplas J, Passos A, Cournapeau D, Brucher M, Perrot M, Duchesnay E (2011) Scikit-learn: Machine Learning in Python. J Mach Learn Res abs/1201.0490 Available at: https://www.jmlr.org/papers/volume12/pedregosa11a/pedregosa11a.pdf.

Prince JS, Charest I, Kurzawski JW, Pyles JA, Tarr MJ, Kay KN (2022) Improving the accuracy of single-trial fMRI response estimates using GLMsingle. Elife 11 Available at: http://dx.doi.org/10.7554/eLife.77599.

Rademaker RL, Chunharas C, Serences JT (2019) Coexisting representations of sensory and mnemonic information in human visual cortex. Nat Neurosci 22:1336–1344.

Reuter M, Rosas HD, Fischl B (2010) Highly accurate inverse consistent registration: a robust approach. Neuroimage 53:1181–1196.

Ricker TJ, Cowan N (2010) Loss of visual working memory within seconds: the combined use of refreshable and non-refreshable features. J Exp Psychol Learn Mem Cogn 36:1355–1368.

Ricker TJ, Vergauwe E, Cowan N (2016) Decay theory of immediate memory: From Brown (1958) to today (2014). Q J Exp Psychol 69:1969–1995.

Rolls ET (2000) Hippocampo-cortical and cortico-cortical backprojections. Hippocampus 10:380–388.

Roth ZN, Heeger DJ, Merriam EP (2018) Stimulus vignetting and orientation selectivity in human visual cortex. Elife 7 Available at: http://dx.doi.org/10.7554/eLife.37241.

Rugg MD, Johnson JD, Park H, Uncapher MR (2008) Encoding-retrieval overlap in human episodic memory: a functional neuroimaging perspective. Prog Brain Res 169:339–352.

Schacter DL, Norman KA, Koutstaal W (1998) The cognitive neuroscience of constructive memory. Annu Rev Psychol 49:289–318.

Schurgin MW, Flombaum JI (2018) Visual working memory is more tolerant than visual long-term memory. J Exp Psychol Hum Percept Perform 44:1216–1227.

Serences JT (2016) Neural mechanisms of information storage in visual short-term memory. Vision Research 128:53–67 Available at: http://dx.doi.org/10.1016/j.visres.2016.09.010.

Serences JT, Ester EF, Vogel EK, Awh E (2009) Stimulus-specific delay activity in human primary visual cortex. Psychol Sci 20:207–214.

Smith AT, Singh KD, Williams AL, Greenlee MW (2001) Estimating receptive field size from fMRI data in human striate and extrastriate visual cortex. Cereb Cortex 11:1182–1190.

Sprague TC, Adam KCS, Foster JJ, Rahmati M, Sutterer DW, Vo VA (2018) Inverted Encoding Models Assay Population-Level Stimulus Representations, Not Single-Unit Neural Tuning. eNeuro 5 Available at: http://dx.doi.org/10.1523/ENEURO.0098-18.2018.

Teyler TJ, Rudy JW (2007) The hippocampal indexing theory and episodic memory: updating the index. Hippocampus 17:1158–1169.

Tong F, Nakayama K, Vaughan JT, Kanwisher N (1998) Binocular rivalry and visual awareness in human extrastriate cortex. Neuron 21:753–759.

Tulving E, Thomson DM (1973) Encoding specificity and retrieval processes in episodic memory. Psychol Rev 80:352–373.

Tustison NJ, Avants BB, Cook PA, Zheng Y, Egan A, Yushkevich PA, Gee JC (2010) N4ITK: improved N3 bias correction. IEEE Trans Med Imaging 29:1310–1320.

Van Essen DC, Maunsell JHR (1983) Hierarchical organization and functional streams in the visual cortex. Trends Neurosci 6:370–375.

Virtanen P et al. (2020) SciPy 1.0: fundamental algorithms for scientific computing in Python. Nat Methods 17:261–272.

Vo VA, Sutterer DW, Foster JJ, Sprague TC, Awh E, Serences JT (2022) Shared Representational Formats for Information Maintained in Working Memory and Information Retrieved from Long-Term Memory. Cereb Cortex 32:1077–1092.

Waskom, M. L., (2021). seaborn: statistical data visualization. Journal of Open Source Software, 6(60), 3021, https://doi.org/10.21105/joss.03021

Wandell BA, Winawer J (2011) Imaging retinotopic maps in the human brain. Vision Res 51:718–737.

Wheeler ME, Petersen SE, Buckner RL (2000) Memory's echo: vivid remembering reactivates sensory-specific cortex. Proc Natl Acad Sci U S A 97:11125–11129.

Winawer J, Witthoft N (2015) Human V4 and ventral occipital retinotopic maps. Vis Neurosci 32:E020.

Xu Y (2017) Reevaluating the Sensory Account of Visual Working Memory Storage. Trends Cogn Sci 21:794–815.

Yassa MA, Stark CEL (2011) Pattern separation in the hippocampus. Trends Neurosci 34:515–525.

Zhang W, Luck SJ (2009) Sudden death and gradual decay in visual working memory. Psychol Sci 20:423–428.

Zhang Y, Brady JM, Smith SM (2001) An hmrf-em algorithm for partial volume segmentation of brain mri fmrib technical report tr01yz1. Brain Available at: https://www.fmrib.ox.ac.uk/datasets/techrep/tr01yz1/tr01yz1.pdf.