

Retrospective Cues Mitigate Information Loss in Human Cortex during Working Memory Storage

 Edward F. Ester,^{1,2,3} Asal Nouri,^{1,3*} and Laura Rodriguez^{1*}

¹Department of Psychology, ²Center for Complex Systems and Brain Sciences, and ³FAU Brain Institute, Florida Atlantic University, Boca Raton, Florida 33431

Working memory (WM) enables the flexible representation of information over short intervals. It is well established that WM performance can be enhanced by a retrospective cue presented during storage, yet the neural mechanisms responsible for this benefit are unclear. Here, we tested several explanations for retrospective cue benefits by quantifying changes in spatial WM representations reconstructed from alpha-band (8–12 Hz) EEG activity recorded from human participants (both sexes) before and after the presentation of a retrospective cue. This allowed us to track cue-related changes in WM representations with high temporal resolution (tens of milliseconds). Participants encoded the locations of two colored discs for subsequent report. During neutral trials, an uninformative cue instructed participants to remember the locations of both discs across a blank delay, and we observed a monotonic decrease in the fidelity of reconstructed spatial WM representations with time. During valid trials, a 100% reliable cue indicated that the color of the disc participants would be probed to report. Critically, valid cues were presented immediately after the termination of the encoding display [“valid early” (VE) trials] or midway through the delay period [“valid late” (VL) trials]. During VE trials, the gradual loss of location-specific information observed during neutral trials was eliminated, while during VL trials it was partially reversed. Our findings suggest that retrospective cues engage several different mechanisms that together serve to mitigate information loss during WM storage.

Key words: alpha band; EEG; retrospective cue; selective attention; working memory

Significance Statement

Working memory (WM) performance can be improved by a cue presented during storage. This effect, termed a retrospective cue benefit, has been used to explore the limitations of attentional prioritization in WM. However, the mechanisms responsible for retrospective cue benefits are unclear. Here we tested several explanations for retrospective cue benefits by examining how they influence WM representations reconstructed from human EEG activity. This approach allowed us to visualize, quantify, and track the effects of retrospective cues with high temporal resolution (on the order of tens of milliseconds). We show that under different circumstances retrospective cues can both eliminate and even partially reverse information loss during WM storage, suggesting that retrospective cue benefits have manifold origins.

Introduction

Visual working memory (WM) enables the representation and manipulation of information that is no longer in the sensorium. This system is an integral component of many higher-order cognitive abilities (Cowan, 2001), yet it has a sharply limited repre-

sentational capacity (Luck and Vogel, 2013; Ma et al., 2014). Thus, mechanisms of selective attention are needed to control access to WM (Vogel et al., 2005) and to prioritize existing WM representations for behavioral output (Myers et al., 2017). Attentional prioritization in WM has been extensively studied using retrospective cues (retro-cue; for recent comprehensive reviews, see Souza and Oberauer, 2016; Myers et al., 2017). In a typical retro-cue experiment, participants encode an array of items for subsequent report. During the ensuing delay period, a cue indicates which of the original items is the most likely to be tested. Relative to a no-cue or neutral-cue baseline, valid cues typically yield greater memory performance while invalid cues typically yield reduced memory performance (although the evidence for invalid cue costs is mixed; Souza and Oberauer, 2016).

Received June 19, 2018; revised Aug. 6, 2018; accepted Aug. 16, 2018.

Author contributions: E.F.E. designed research; A.N. and L.R. performed research; E.F.E., A.N., and L.R. analyzed data; E.F.E. wrote the paper.

This work was supported by a State University System of Florida startup grant to E.F.E. We thank Tommy Sprague for helpful comments on an earlier version of this manuscript.

*A.N. and L.R. contributed equally to the project and are listed alphabetically.

The authors declare no competing financial interests.

Correspondence should be addressed to Edward F. Ester, Department of Psychology, Florida Atlantic University, BS-12, Room-101, 777 Glades Road, Boca Raton, FL 33431. E-mail: eester@fau.edu.

DOI:10.1523/JNEUROSCI.1566-18.2018

Copyright © 2018 the authors 0270-6474/18/388538-11\$15.00/0




Figure 1. Spatial memory task. Participants encoded the locations of two colored discs (red and blue; sample display). Each disc was presented at one of nine equally spaced locations along the perimeter of an imaginary circle centered at fixation (see Materials and Methods; Fig. 3). During VE trials, the color of the fixation point changed from black to either blue or red immediately after termination of the sample display. This cue indicated which disc would be probed with 100% reliability. During neutral trials, the color of the fixation point changed from black to purple, instructing participants to remember the locations of both discs. During VL trials, the color of the fixation point changed from black to purple. Halfway through the subsequent delay period, the fixation point changed colors from purple to either blue or red. This second change indicated which disc would be probed with 100% reliability. Each trial concluded with a probe display containing a red or blue fixation point, a question mark, and a mouse cursor. Participants were instructed to report the precise location of the disc indicated by the color of the fixation point via mouse click. Note: the schematic is included for illustrative purposes; displays are not drawn to scale. See Materials and Methods for display and stimulus geometry.

Several mechanisms may be responsible for retrospective cue benefits in WM performance. For example, multiple studies have reported reductions in load-dependent neural signatures of WM storage following a retrospective cue, suggesting that these cues engage mechanisms that facilitate the removal of irrelevant items from WM (Williams and Woodman, 2012; Kuo et al., 2012). Other studies have reported changes in lateralized alpha-band activity when participants are retrospectively cued to an item that previously appeared in the left or right visual hemifields, which is consistent with an attentional prioritization of the cued representation (Poch et al., 2014; Myers et al., 2015; van Moorselaar et al., 2018) and/or suppression of the uncued representation. There is also evidence suggesting that valid retrospective cues trigger mechanisms that insulate cued WM representations from subsequent display or interference (Makovski et al., 2010; Pertzov et al., 2013), or mechanisms that facilitate comparisons between cued WM representations and subsequent memory probes (Souza and Oberauer, 2016).

Here, we tested several explanations for retro-cue benefits by quantifying changes in spatially specific WM representations before and after the appearance of a retrospective cue. Inspired by earlier work (Foster et al., 2016), we reconstructed spatially specific mnemonic representations by applying an inverted encoding model (IEM) to spatiotemporal patterns of alpha-band (8–12 Hz) activity recorded while participants performed a retrospectively cued spatial WM task. On each trial, participants encoded the locations of two colored discs (blue and red). During neutral trials, an uninformative color cue presented after the encoding display informed participants to remember the locations of both discs across a blank interval. During valid trials, a 100% reliable color cue indicated which disc would be probed at the end of the trial. Valid cues were presented either immediately after offset of the encoding display [“valid early” (VE) trials] or at the midpoint of the subsequent blank period [“valid late” (VL) trials]. We isolated the effects of retrospective cues on spatial WM performance

by comparing location-specific WM representations across the neutral and valid conditions. To preview the results, during neutral trials we observed a gradual decrease in the fidelity of location-specific representations with time. During VE trials, this decrease was eliminated, while during VL trials it was partially reversed. Our findings thus support the view that retro-cues can engage several different mechanisms that together serve to mitigate information loss during WM storage.

Materials and Methods

Participants. Thirty-one volunteers from the Florida Atlantic University community (age, 18–40 years; both sexes) completed a single 2.5 h testing session. All participants self-reported normal or corrected-to-normal visual acuity and were compensated at a rate of \$15/h. Data from four participants were discarded due to an excessive number of electro-oculogram (EOG) artifacts (>33% of trials). Data from a fifth participant was discarded as that participant withdrew from the study after completing only 6 of 12 testing blocks. The data reported here reflect the remaining 26 volunteers.

Testing environment. Participants were seated in a dimly lit and sound-attenuated (unshielded) recording chamber. Stimuli were generated in MATLAB using Psychtoolbox-3 software (all data and analytical software associated with this project are publicly available on the Open Sciences Framework: <https://osf.io/m4c2p/>; Kleiner et al., 2007) and rendered on a 17-inch Dell CRT monitor cycling at 85 Hz. Participants were seated ~60 cm from the display (head position was not constrained).

Spatial memory task. The behavioral task was based on an experimental approach described by Sprague et al. (2016). A trial schematic is shown in Figure 1. Participants were instructed to fixate a small dot (subtending 0.2° from a viewing distance of 60 cm) throughout the experiment. Each trial began with a sample display containing two colored discs (one red and one blue). Each disc was presented in one of nine equally spaced positions (from 40° to 360° in 40° increments) along the perimeter of an imaginary circle (radius 6° visual angle) centered at the fixation point. A small amount of jitter ($\pm 10^\circ$ polar angle) was added to the location of each disc on each trial to discourage verbal coding strategies (e.g., “the blue disc was at 2 o’clock”).

The sample display was extinguished after 500 ms. During VE trials, the fixation point changed colors from black to either blue or red. This change served as a 100% valid cue for the disc whose location participants would be asked to report at the end of the trial. The fixation point remained blue (or red) for the remainder of the delay period. During VL trials, the color of the fixation point initially changed from black to purple, instructing participants to remember the locations of both discs. At the midpoint of the delay period (1250 ms after offset of the sample display), the fixation point changed colors from purple to either blue or red; this change served as a 100% valid cue for the disc whose location participants would be asked to report at the end of the trial. The fixation dot remained red or blue for the remainder of the second delay period (1251–2500 ms after sample offset). Finally, during neutral trials the color of the fixation point changed from black to purple and remained purple across the entire 2500 ms delay period.

Each trial concluded with a test display containing a blue or red fixation point, a mouse cursor, and a question mark symbol (“?”) above the fixation point. Participants were required to click on the location of the disc indicated by the color of the fixation point within a 3000 ms response window. Memory performance was quantified as the absolute angular distance between the polar location of the probed disc and the polar

location reported by the participant. Performance feedback (mean absolute recall error; i.e., the mean absolute difference between the polar angle reported by the participant and the polar location of the probed disc) was given at the end of each block. Participants completed 10 ($N = 1$), 11 ($N = 2$), or 12 ($N = 23$) blocks of 72 trials. Cue conditions (VE, VL, and neutral) were counterbalanced within each block, while the spatial positions of the red and blue discs were counterbalanced across sessions of six blocks.

EEG acquisition and preprocessing. Continuous EEG was recorded from 62 scalp electrodes via a BrainProducts actiCHamp Amplifier. Additional electrodes were placed over the left and right mastoids. Electrode impedances were kept to <15 k Ω . Data were recorded with a right mastoid reference and digitized at 1000 Hz. The horizontal EOG was recorded from electrodes located ~ 1 cm from the left and right canthi, while the vertical EOG was recorded from electrodes placed above and below the right eye. Data were later rereferenced to the algebraic mean of the left and right mastoids, bandpass filtered from 0.1 to 40 Hz (third-order zero-phase forward-and-reverse Butterworth filters), epoched from -1000 to $+4000$ ms relative to the start of each trial, and baseline corrected from -100 to 0 ms relative to the start of each trial. Trials contaminated by blinks or horizontal eye movements greater than $\sim 2.5^\circ$ (assuming a normative voltage threshold of $16 \mu\text{V}^\circ$; Lins et al., 1993) were excluded from subsequent analyses. Noisy electrodes were identified and removed by visual inspection. Across participants, we rejected an average (± 1 SEM) of 12.38% ($\pm 1.83\%$) trials and 1.04 (± 0.23) electrodes.

Inverted encoding model. Following earlier work (Foster et al., 2016), we used an IEM to reconstruct location-specific representations of the red and blue discs during the sample display and subsequent delay period. Reconstructions of stimulus locations were computed from the spatial topography of induced alpha-band (8–12 Hz) power measured across the following 17 occipitoparietal electrode sites: O1, O2, Oz, PO7, PO3, POz, PO4, PO8, P7, P5, P3, P1, Pz, P2, P4, P6, and P8. Robust representations of spatial position could also be reconstructed from frontal, central, and temporal electrode sites, but these representations were substantially weaker than those reconstructed from occipitoparietal electrode sites. To isolate alpha-band activity, the raw EEG time series at each electrode was bandpass filtered from 8 to 12 Hz (third-order zero-phase forward-and-reverse Butterworth filter), yielding a real-value signal, $f(t)$. The analytic representation of $f(t)$ was obtained by applying a Hilbert transformation, as follows:

$$z(t) = f(t) + if(t),$$

where $i = \sqrt{-1}$ and if $(t) = A(t)e^{i\varphi(t)}$. Induced alpha power was computed by extracting and squaring the instantaneous amplitude $A(t)$ of the analytic signal $z(t)$.

We modeled alpha power at each scalp electrode as a weighted sum of nine location-selective channels, each with an idealized tuning curve (a half-wave rectified cosine raised to the eighth power). The maximum response of each channel was normalized to 1; thus, the units of response are arbitrary. The predicted responses of each channel during each trial were arranged in a k channel by n trials design matrix C . Separate design matrices were constructed to track the locations of the blue and red discs across a trial (i.e., we reconstructed the locations of the blue and red discs separately, then later sorted these reconstructions according to cue condition).

The relationship between the data and the predicted channel responses C is given by a general linear model of the following form:

$$B = WC + N,$$

where B is an m electrode by n trials training data matrix, W is an m electrode by k channel weight matrix, and N is a matrix of residuals (i.e., noise).

To estimate W , we constructed a “training” dataset containing an equal number of trials from each stimulus location (i.e., $40\text{--}360^\circ$ in 40° steps). We first identified the location φ with the fewest r repetitions in the full dataset after EOG artifact removal. Next, we constructed a training (trn) dataset B_{trn} (m electrodes by n trials) and weight matrix C_{trn}

(n trials by k channels) by randomly selecting (without replacement) $1:r$ trials for each of the nine possible stimulus locations (ignoring cue condition; i.e., the training dataset contained a mixture of VE, VL, and neutral trials). The training dataset was used to compute a weight for each channel, C_p , via least-squares estimation, as follows:

$$W_i = B_{\text{trn}} C_{\text{trn},i}^T (C_{\text{trn},i} C_{\text{trn},i}^T)^{-1},$$

where $C_{\text{trn},i}$ is an n trial row vector containing the predicted responses of spatial channel i during each training trial.

The weights W were used to estimate a set of spatial filters, V , that capture the underlying channel responses while accounting for correlated variability between electrode sites (i.e., the noise covariance; Kok et al., 2017), as follows:

$$V_i = \frac{\sum_i^{-1} W_i}{W_i^T \sum_i^{-1} W_i},$$

where Σ_i is the regularized noise covariance matrix for channel i and is estimated as follows:

$$\Sigma_i = \frac{1}{n-1} \epsilon_i \epsilon_i^T,$$

where n is the number of training trials and ϵ_i is a matrix of residuals, as follows:

$$\epsilon_i = B_{\text{trn}} - W_i C_{\text{trn},i}.$$

Estimates of ϵ_i were obtained by regularization-based shrinkage using an analytically determined shrinkage parameter (Blankertz et al., 2011; Kok et al., 2017). An optimal spatial filter, v_p , was estimated for each channel C_p , yielding an m electrode by k filter matrix V .

Next, we constructed a “test” (tst) dataset, B_{tst} (m electrodes $\times n$ trials) containing data from all trials not included in the training dataset and estimated trial-by-trial channel responses C_{tst} (k channels $\times n$ trials) from the filter matrix V and the test dataset, as follows:

$$C_{\text{tst}} = V^T B_{\text{tst}}.$$

Trial-by-trial channel responses were interpolated to 360° , circularly shifted to a common center (0° , by convention), and sorted by cue condition (i.e., VE, VL, and neutral). To quantify the effects of retrospective cues on spatial memory representations, we obtained an estimate of location-specific information by converting the (averaged) channel response estimates for each cue condition to polar form and projected them onto a vector with angle 0° , as follows:

$$r = |z| \cos(\arg(z)), z = ce^{2i\varphi},$$

where c is a vector of estimated channel responses and φ is the vector of angles at which the channels peak.

To ensure internal reliability, this entire analysis was repeated 100 times, and unique (randomly chosen) subsets of trials were used to define the training and test datasets during each permutation. The results were then averaged across permutations.

Estimating the temporal resolution of location information estimates. As described in the preceding section, location-specific reconstructions (and hence, estimates of location information) were computed using instantaneous distributions of alpha power across electrode sites at each time point. However, the bandpass filter used to isolate alpha-band activity (8–12 Hz) can introduce temporal distortions, making it difficult to make precise statements about the temporal resolution of the analysis pipeline. To formally investigate this issue, we generated a 3000 ms sinusoid with unit amplitude and a frequency of 10 Hz (along with 1000 ms of presignal and postsignal zero padding for a total signal length of 5000 ms) and ran it through the filtering routine used to identify alpha-band activity (defined as 8–12 Hz) in the raw EEG signal. Temporal distortions caused by filtering were defined as points where the amplitude of the filtered signal reached 25% of maximum (1.0). For a perfect filter, these points would occur at precisely 1000 and 4000 ms relative to stimulus

onset. In reality, we obtained estimates of 963 and 4039 ms, or 37 ms before and 39 ms after stimulus onset. Thus, we estimated the temporal resolution of our analysis path as approximately ± 40 ms. That is, if estimates of location information are significantly above zero at a given time point, then we can be certain that the spatial distribution of alpha power across electrode sites contained robust information about that location at some point within ± 40 ms.

Statistical comparisons. We used a nonparametric sign permutation test to quantify differences between estimates of location information across stimuli (cued vs uncued) and cue conditions (i.e., VE, VL, and neutral trials). Each test we performed (e.g., whether estimated location information is reliably >0 or whether estimated location information is higher for the cued vs uncued disc) assumes a null statistic of 0. We therefore generated a null distribution of location information estimates by randomly inverting the sign of each participant's data (with 50% probability) and averaging the data across participants. This procedure was repeated 10,000 times, yielding a 10,000-element null distribution for each time point. Finally, we implemented a cluster-based permutation test with cluster-forming and cluster-size thresholds of $p < 0.05$ to correct for multiple comparisons across time points (Maris and Oostenveld, 2007; Wolff et al., 2017).

Eye movement control analyses. Although we excluded trials contaminated by horizontal eye movements, small but reliable biases in eye position toward the locations of the remembered discs could nevertheless contribute to reconstructions of stimulus location. We examined this possibility in two complementary analyses. In the first analysis, we computed angular measurements of eye position by converting trial-by-trial horizontal and vertical EOG recordings to polar format (Lins et al., 1993; Bae and Luck, 2018; using normative scaling values of 16 and 12 μV° for the horizontal and vertical EOG channels, respectively), then constructed circular histograms of angular estimates of eye position as a function of stimulus location. Angular eye position estimates at each time point were sorted into nine bins whose centers matched the possible stimulus positions (i.e., from 40° to 360° in 40° increments). For simplicity, we restricted our analysis to VE trials where consistent biases in eye position should be most apparent (i.e., because participants were only required to remember one location).

In the second analysis, we regressed trial-by-trial horizontal EOG recordings (in microvolts; Foster et al., 2016) onto the horizontal position of the remembered disc, again focusing on VE trials. Positive regression coefficients thus reflect greater changes in eye position as a function of stimulus location. Separate regressions were run for each participant, and the resulting coefficients were averaged across participants.

Quantifying sources of location information loss and recovery. We evaluated potential sources of change in location-specific representations via curve-fitting analyses (Ester et al., 2015, 2016). We first computed a one-dimensional reconstructed representation of location for each participant by averaging channel responses over time (separately for the first and second delay periods). Reconstructions were averaged across participants and fit with a circular function of the form, as follows:

$$f(x) = \alpha e^{k \cos(\mu - x) - 1} + \beta,$$

where x is a vector of polar locations and β is an additive scaling parameter. The additional parameters α , μ , and k control the amplitude, center, and concentration (i.e., the inverse of bandwidth) of the function. α , β , μ , and k were estimated separately for each condition via an iterative minimization algorithm (Nelder–Mead simplex as implemented by the MATLAB “fminsearch” function).

Differences in parameters across conditions (i.e., first vs second delay and neutral vs VE vs VL) were evaluated using a bootstrap test. For each comparison, we selected (with replacement) and averaged 26 of 26 participant reconstructions from the two conditions of interest (e.g., neutral delay 1 and VE delay 1). The resulting average functions were fit with the circular function described above, yielding a set of parameter estimates for each condition. This procedure was repeated 10,000 times, and we computed empirical p values for differences in parameter estimates across conditions by estimating the total proportion of permutations where parameter estimates in one condition were greater than (or less than) estimates for the other condition.




Figure 2. Behavioral performance. Recall error (i.e., the absolute polar angle between the location reported by the participant and the location of the probed disc) was reliably lower during VE trials relative to neutral and VL trials. Recall error was also reliably lower during VL relative to neutral trials. Error bars show 95% confidence intervals.

Within-participant variability. We report estimates of within-participant variability (e.g., 95% within-participant confidence intervals) throughout the article. These estimates discard subject variance (e.g., overall differences in response strength) and instead reflect variance related to the subject \times conditions interaction terms (e.g., variability in response strength across experimental conditions; Loftus and Masson, 1994; Cousineau, 2005). We used the approach described by Cousineau (2005), as follows: raw data (e.g., location information or channel response estimates) were demeaned on a participant-by-participant basis, and the grand mean across participants was added to each participant's zero-centered data. The grand mean-centered data were then used to compute bootstrapped within-participant confidence intervals (10,000 permutations).

Results

Valid retro-cues enhance memory performance

We recorded EEG while participants performed a retrospectively cued spatial WM task (Fig. 1). During neutral trials, an uninformative cue instructed participants to remember the locations of two colored discs across a blank delay period. During valid trials, a 100% reliable color cue indicated which disc would be probed at the end of the trial. Valid cues were presented immediately after termination of the encoding display (VE trials) or midway through the subsequent blank delay period (VL trials). At the end of the trial, participants recalled the location of the probed disc via a mouse click. Behavioral performance was quantified as the mean absolute recall error (i.e., the mean absolute difference between the polar angle reported by the participant and the polar location of the probed disc). Figure 2 plots behavioral performance as a function of trial type (i.e., neutral, VE, and VL). Consistent with earlier findings (Sprague et al., 2016), the average absolute recall error was reliably lower during VE and VL trials (mean, 7.75° and 8.26° , respectively) relative to neutral trials (mean, 8.96° ; false discovery rate-corrected bootstrap tests, $p < 0.001$). The average absolute recall error was also reliably lower during VE relative to VL trials (bootstrap test, $p < 0.001$). Thus, valid retro-cues improved WM performance.

Reconstructing location-specific WM representations from alpha-band activity

Building upon earlier work (Foster et al., 2016; Samaha et al., 2016), we used an IEM to reconstruct the locations of the cued and uncued discs from spatiotemporal patterns of induced alpha-band power measured across occipitoparietal electrode sites (see Inverted encoding model, in Materials and Methods). Briefly, we modeled instantaneous induced alpha power at each electrode site as a weighted combination of nine location-selective chan-

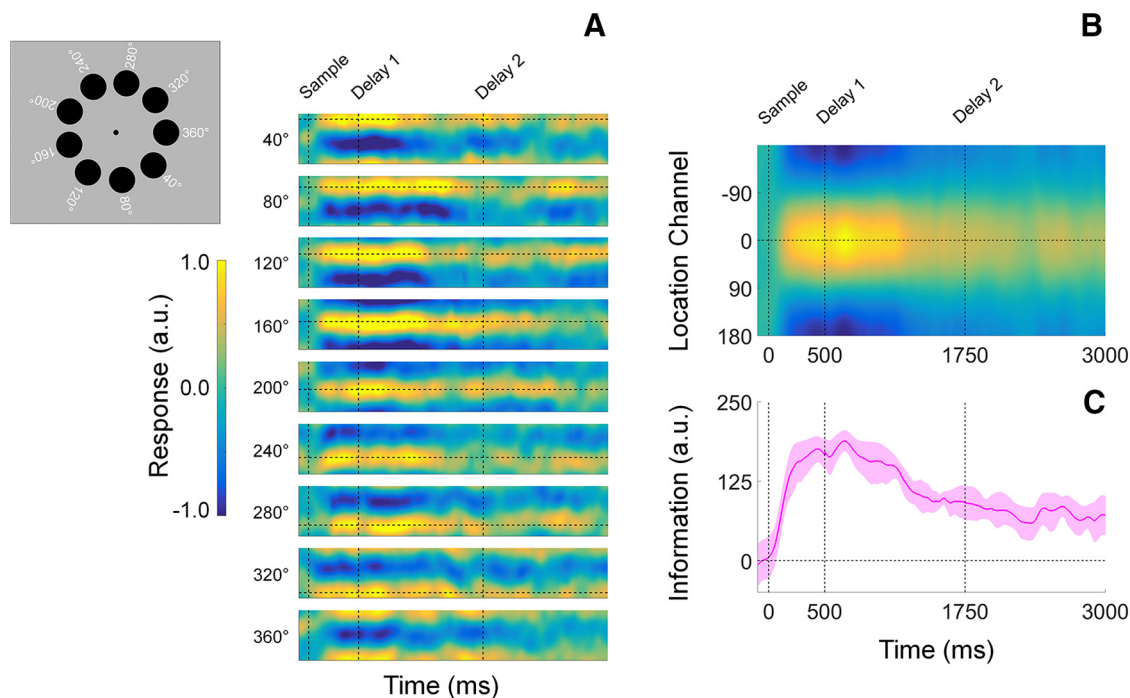


Figure 3. Computing time-resolved reconstructions of stimulus location. **A**, Location reconstructions for each of the nine possible stimulus locations from 40° to 360°. The dashed horizontal line in each plot shows the polar location of the stimulus, while the vertical dashed lines at times 0, 500, and 1750 mark the start of the sample epoch, first delay period, and second delay period, respectively (the time axis is identical to that shown in **B** and **C**). The inset to the left shows each of the nine possible disc locations on each trial. Reconstructions have been pooled and averaged across stimulus identity (i.e., red vs blue disc) and cue condition (VE, VL, and neutral). **B**, We circularly shifted the reconstructions shown in **A** to a common center (0°) and averaged them, yielding a single time-resolved location reconstruction. The response scale is identical to that shown in **A**. **C**, We converted the reconstructions shown in **B** to polar form and projected them onto a vector with angle 0°. We interpreted the resultant vector length as a measure of total location-specific information. Shaded regions represent the 95% within-participant confidence interval.

nels, each with an idealized response function. The resulting channel weights were then used to calculate a predicted response for each channel, given the spatiotemporal distributions of alpha power measured during an independent test dataset. Separate reconstructions were computed for each disc regardless of cue condition, and the individual reconstructions were averaged to yield a single time-resolved representation of location-specific activity. Consistent with earlier findings (Foster et al., 2016) trial-by-trial variability in channel responses reliably tracked the angular locations of the cued and uncued discs (Fig. 3A). For convenience, we circularly shifted reconstructions for each stimulus location to a common center (0° by convention) and averaged the centered reconstructions across locations, yielding a single time-resolved reconstruction. Finally, we converted these centered reconstructions to polar form and projected them onto a unit vector with an angle of 0°. As shown in Figure 3C, location information increased rapidly following presentation of the sample display and reached an asymptotic limit ~400–500 ms later. During the subsequent delay period, location information gradually decreased with time, although overall information levels remained reliably greater than 0 for the duration of the trial. Thus, our analytical approach allowed us to visualize spatially specific WM representations with high temporal precision.

Comparison of spatial representations during encoding

The central goal of this study was to examine the effects of retrospective cues on reconstructed spatial WM representations. Although participants had no way of knowing what type of cue would be presented on each trial, it is conceivable that differences in reconstructions across cue conditions during the encoding phase of the trial contributed to subsequent differences across

retrospective cue conditions. We evaluated this possibility by directly comparing estimates of location information across retrospective cue conditions (VE, VL, and neutral; Fig. 4A). As shown in Figure 4B, estimates of location information were remarkably consistent across cue conditions, and we failed to identify any statistically robust differences in location information as a function of cue condition over the entire 500 ms encoding period (false discovery rate-corrected, cluster-based permutation tests, $p > 0.05$). Thus, we can be certain that any effects of cue type during the subsequent WM period are not due to differences that emerged during encoding.

Degradation of spatial WM representations during neutral trials

Next, we examined the effects of retro-cues on reconstructed spatial WM representations. Since all retrospective cues were presented after the offset of the sample display, we limited our analyses to the 2500 ms blank interval separating the offset of the sample display and the onset of the probe display.

During neutral trials, a retrospective cue presented immediately after offset of the sample display instructed participants to remember the locations of both discs. As shown in Figure 5A, location information decreased monotonically over the course of the delay period (linear slope = -62.97 units/s; bootstrap test, $p < 0.002$) with information reaching levels indistinguishable from 0 by the onset of the probe display. We next examined possible sources of information loss, including reductions in reconstruction amplitude (i.e., a lower signal-to-noise ratio; Sprague et al., 2014, 2016) or increases in reconstruction bandwidth (i.e., a loss of precision; Ester et al., 2013). We averaged each participant's time-resolved location reconstructions over

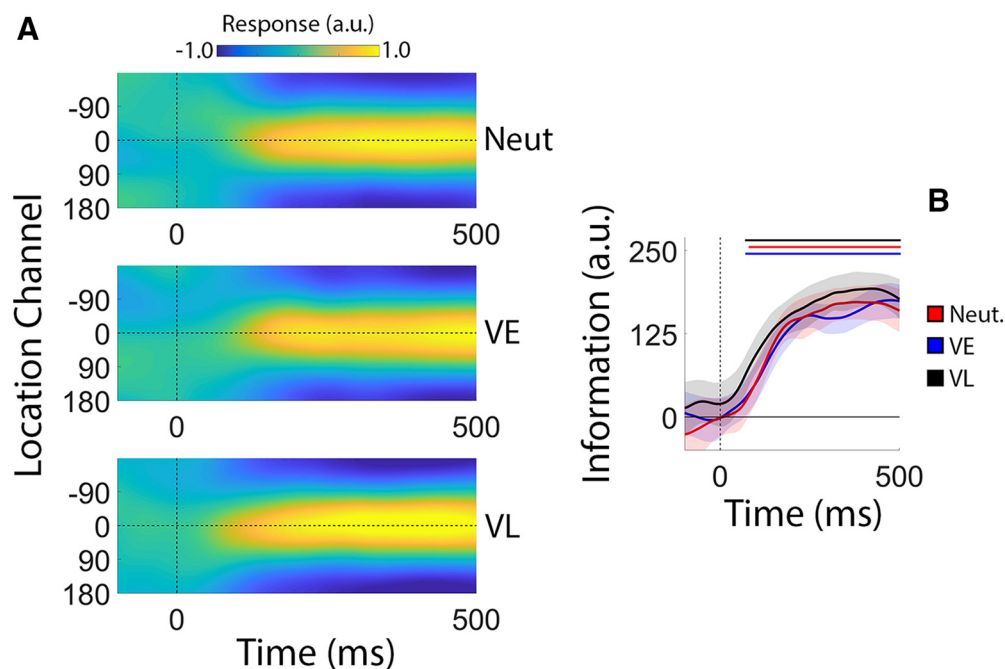


Figure 4. Stimulus reconstructions and location information during encoding. Participants encoded the locations of two colored discs for subsequent report (Fig. 1). **A**, Channel response functions were identical across cue conditions. This is unsurprising, as participants had no way of knowing what type of cue they would receive until termination of the encoding display. **B**, Robust location information emerged ~ 100 ms after display onset and increased rapidly before reaching an asymptotic limit ~ 350 ms after display onset. Statistically, estimates of location information were identical across cue conditions. Thus, we can be certain that differences in location information during the subsequent WM period are not due to differences that emerge during encoding. Solid lines at the top of **B** mark epochs where estimates of location information were reliably greater than 0. Shaded regions depict the 95% within-participant confidence interval of the mean.




Figure 5. Degradation of spatial WM representations during neutral trials. **A**, During neutral trials, we observed a monotonic decrease in total location information during the delay period. Data have been pooled and averaged across stimulus locations (i.e., the locations of the blue and red discs) and are time locked to the offset of the sample display (0 ms). The pink bar at the top of the plot marks epochs where estimates of location information were reliably greater than 0 (false discovery rate-corrected, cluster-based permutation test; see Materials and Methods). **B**, We identified sources of information loss by computing and quantifying time-averaged, location-specific reconstructions during the first and second delay periods (see Quantifying sources of information loss and recovery, in Materials and Methods). **C**, Reconstruction amplitudes were reliably lower during the second delay period relative to the first delay period, suggesting that information loss reflects a gradual reduction in the overall strength of each spatial WM representation. For all plots, shaded regions and error bars show 95% within-participant confidence intervals.

periods from 0 to 1250 and 1251 to 2500 ms after the offset of the sample display (Fig. 5B). Each reconstruction was fit with a circular function containing free parameters for amplitude (i.e., maximum over baseline response), concentration (the inverse of bandwidth), and center (see Quantifying sources of information loss and recovery, in Materials and Methods). As shown in Figure 5C, we observed a decrease in response amplitude across the first and second halves of the delay period, but no change in concentration. Thus, requiring participants to store multiple spatial WM representations was associated with a gradual decrease in the strength of each representation over time.

Valid retro-cues presented immediately after encoding eliminate information loss

During VE trials, a retrospective cue presented immediately after the sample display indicated which disc (blue or red) would be

probed with 100% certainty. In contrast to the pattern seen during neutral trials (Fig. 5), location information for the cued disc remained constant across the delay period (Fig. 6; linear slope, -17.8 units/s; bootstrap test, $p = 0.06$). Contrary to prior results suggesting a cue-driven increase in information over and above that observed during encoding (Rerko et al., 2014; Schneegans and Bays, 2017), we found no evidence for a strengthening of the cued location immediately after cue onset (obtained by comparing the average location information over epochs spanning the last 100 ms of the sample period and the first 500 ms of the delay period; bootstrap test, $p = 0.93$). Conversely, location information for the uncued disc quickly fell to 0. The rate of information loss for the uncued disc during the first 1250 ms of the delay period during VE trials was nearly double that observed during the same interval of neutral trials (linear slopes of -143.61 vs

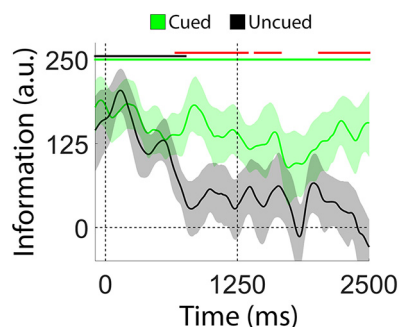


Figure 6. Information loss is prevented by a retrospective cue presented immediately after encoding. During VE trials, a 100% reliable retrospective cue indicated which disc (blue or red) participants would be asked to report. During these trials, we observed a rapid decrease in location information for the uncued disc (black line), but no change in location information for the cued disc (green line). Green and black bars at the top of the plot mark epochs where location information was reliably greater than 0 for the cued and uncued discs, respectively. Red bars mark epochs where location information was reliably larger for the cued relative to the uncued disc. Shaded regions depict 95% within-participant confidence intervals.

74.18 units/s; bootstrap test, $p < 0.02$). Clear differences between location information estimates for the cued and uncued discs emerged ~ 600 ms after cue onset and persisted throughout the remainder of the delay period (Fig. 6B, red markers). This delay is consistent with behavioral studies suggesting that it takes participants ~ 300 – 600 ms to process and use a retrospective cue (Souza et al., 2014).

To summarize, we observed a monotonic decrease in location information with time when participants were required to hold two locations in WM (Fig. 5). A retrospective cue presented immediately after termination of the sample display eliminated this decline for the cued location and hastened it for the uncued location (Fig. 6). The observation that a retrospective cue mitigates information loss for a cued item is consistent with behavioral and physiological findings suggesting that retrospectively cued shifts of attention insulate cued WM representations from subsequent degradation due to interference or decay (Pertsov et al., 2013).

Recovery of location information following a delayed retrospective cue

In a recent study, Sprague et al. (2016; see also Rose et al., 2016; Wolff et al., 2017) documented an apparent recovery of location-specific information following a delayed retro-cue. We tested for a similar effect by examining the effect of a delayed retro-cue on location information (Fig. 7). During VL trials, a neutral cue instructed participants to remember the locations of both discs. Halfway through the delay period, the neutral cue changed colors to either blue or red, indicating with 100% probability which disc would be probed. As shown in Figure 7A, location information decreased gradually over the course of the first delay period (0–1250 ms after the sample display; linear slope, -120.78 units/s; bootstrap test, $p < 0.002$) and continued to decline for the uncued disc during the second delay period (Fig. 7A, black line). Conversely, location information for the cued item appeared to increase during the second delay period (Fig. 7A, green line). To evaluate these changes, we computed the average location information for the cued and uncued discs after dividing the second delay period into early and late epochs (1250–1700 and 2050–2500 ms after sample offset, respectively; Fig. 7B) based on visual inspection of the plots shown in Figure 7A. Estimates of location information were identical for the cued and uncued discs during the early epoch and diverged during the late epoch. Direct com-

parisons of information estimates during the early and late epochs revealed reliably larger information estimates for the cued disc ($p < 0.007$) and reliably lower information estimates for the uncued disc ($p < 0.015$) during the late epoch. We examined sources of location information recovery via a curve-fitting analysis. Specifically, we computed and quantified the amplitude and concentration of time-averaged reconstructions of the cued disc during the early and late portions of the second delay period (i.e., 1250–1700 and 2050–2500 ms; Fig. 7C). Consistent with earlier findings (Sprague et al., 2016), we observed significantly larger amplitude estimates during the late relative to the earlier epoch (Fig. 7D; $p < 0.001$; the difference between concentration estimates during the early and late periods was not significant, $p = 0.0543$). Thus, in addition to protecting representations of cued items from subsequent degradation (Fig. 5), under some circumstances a retrospectively cued shift of attention can directly enhance the representation of a cued item (Fig. 7).

Ruling out contributions from eye movements

We identified and excluded trials contaminated by horizontal eye movements greater than 2.5° (based on a normative voltage threshold of $16 \mu\text{V}^\circ$; Lins et al., 1993). Nevertheless, small but consistent biases in eye position may have contributed to reconstructions of stimulus location. We examined this possibility in two complementary analyses (see Eye movement control analyses, in Materials and Methods). First, we computed angular estimates of eye position from horizontal and vertical EOG recordings and constructed circular histograms of eye position as a function of remembered stimulus location (Fig. 8A). Distributions of eye position were remarkably similar across stimulus locations, suggesting that the location-specific reconstructions shown in Figures 3, 4, 5, 6, and 7 cannot be explained solely by subtle biases in eye position. This conclusion was further supported by a complementary analysis in which we regressed time-resolved estimates of horizontal EOG activity onto remembered stimulus locations (Fig. 8B). Time-resolved regression coefficients were indistinguishable from 0 across each trial, again suggesting that systematic biases in eye position are insufficient to account for the location-specific reconstructions reported here.

Discussion

WM performance can be improved by a retrospective cue presented after encoding is complete (Griffin and Nobre, 2003; Landman et al., 2003; Souza and Oberauer, 2016; Myers et al., 2017). Several mechanisms have been proposed to explain retrospective cue benefits in WM performance, including the removal of irrelevant information from WM (Kuo et al., 2012; Souza and Oberauer, 2016), attentional enhancement of the cued representation (Myers et al., 2015), protection of the cued representation from subsequent decay or interference (Makovsik and Jiang, 2007; Pertsov et al., 2013; e.g., due to competition from other memory representations or subsequent sensory input), or retrieval head start (Souza et al., 2016). Evaluating these alternatives has proven difficult, in part because extant studies examining the effects of retrospective cues have relied almost exclusively on behavioral reports and/or indirect neural signatures of WM storage. In the current study, we overcame this limitation by directly examining changes in spatially specific WM representations before and after the appearance of a retrospective cue. Our approach was predicated on recent studies demonstrating that topographic distributions of induced alpha-band activity encode precise location-specific information during covert attention tasks (Foster et al., 2017) and spatial WM tasks (Foster et al.,

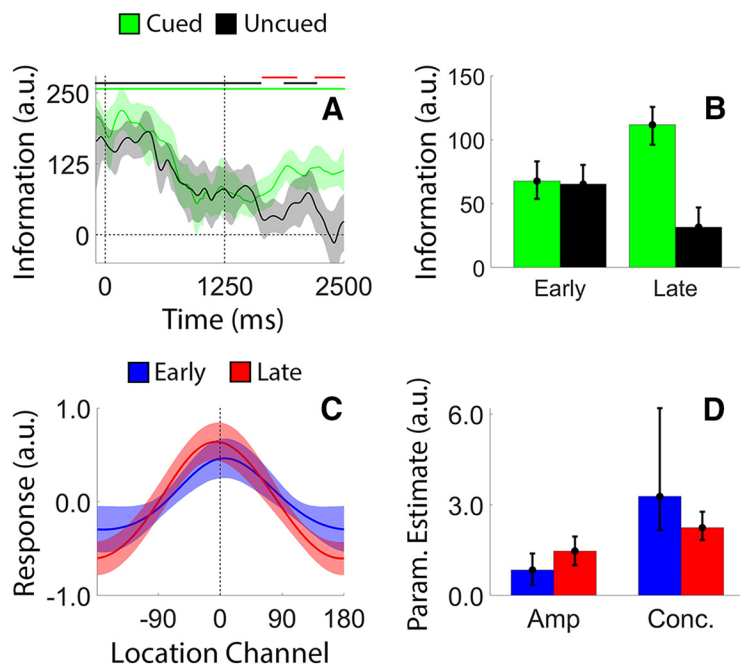


Figure 7. Recovery of location information following a delayed retrospective cue. During VL trials, a neutral cue was replaced by a valid cue midway through the delay period. **A**, We observed a monotonic decrease in location information during the neutral portion of the trial, followed by a partial recovery of location information for the cued disc during the valid portion of the trial. Pink, green, and black lines at the top of the plot mark epochs where estimates of location information were reliably >0 , while red lines mark epochs where the estimates of location information were reliably greater for the cued disc relative to the uncued disc. **B**, To quantify changes in location information after retro-cue onset, we divided the second delay period into separate early and late epochs (1251–1700 and 2050–2500 ms after sample offset, respectively) and computed the average location information for each disc across both epochs. Location information was significantly greater for the cued disc during the late relative to the early epoch, while location information was significantly smaller for the cued disc during the late relative to the early epoch. **C**, We computed and quantified time-averaged reconstructions of the location of the cued disc during the early and late epochs of the second delay period. **D**, Reconstruction amplitudes were reliably larger during the late relative to the early epoch, which is consistent with an attention-based enhancement of the cued representation. For all plots, shaded regions and error bars depict 95% within-participant confidence intervals.

2016; van Ede, 2017), and it allowed us to visualize and quantify cue-driven changes in WM representations with high temporal resolution. Our key findings are summarized in Figure 9. During neutral trials, an uninformative retrospective cue instructed participants to remember the locations of two discs across a blank display, and we observed a gradual degradation in the strength of location-specific WM representations with time (Fig. 4). During VE trials, a 100% reliable retrospective cue was presented immediately after termination of the sample display. This cue eliminated time-based degradation for the representation of the cued disc and hastened it for the uncued disc (Fig. 5). Finally, during VL trials a neutral retrospective cue was replaced with a 100% reliable retrospective cue midway through the delay period. We observed a gradual degradation in location-specific WM representations during the neutral phase of the trial followed by a partial recovery of location during the valid phase of the trial (Fig. 6). Collectively, our findings suggest that retrospective cues can engage multiple mechanisms to minimize and/or reverse information loss during WM storage (Sprague et al., 2015).

The loss of location information observed during neutral and VL trials need not imply a loss of memory. Indeed, the fact that participants perform quite well during neutral trials, despite an apparent absence of location information, suggests at least a partial dissociation between alpha-band activity and spatial WM performance. However, there are at least two ways to explain this. One possibility is that spatial WM representations remain stable

during the delay period, but location information carried by the alpha-band signal degrades with time. A second possibility is that the memory representations are (partially or wholly) consolidated into a new format not indexed by alpha-band activity. The pattern of results we observed during VL trials supports the latter alternative. Under the constraints of information theory (specifically, the data-processing inequality theorem; Shannon, 1948), the total information about one variable given the state of another variable (i.e., mutual information) cannot be increased through additional processing. For example, applying multiplicative gain to a noisy spatial WM representation (e.g., by shifting attention to the cued location) would amplify signal and noise to the same extent, resulting in a stronger response but no increase in the information content of the signal (Sprague et al., 2016; Bays and Taylor, 2018). Thus, once location information is lost it cannot be recovered through any amount of additional processing, unless participants have access to an additional source of information. Since there was no external source of location information during the delay period, we can infer that participants were able to access an internal source of information that is not indexed by alpha-band activity, including but not limited to an “activity-silent” WM system (Rose et al., 2016; Sprague et al., 2016; Wolff et al., 2017) or long-term memory (Sutterer et al., 2018).

Psychophysical studies suggest that retrospective cues engage mechanisms that insulate or protect cued representations from subsequent degradation (Matsukura et al., 2007; Pertzov et al., 2013). Our findings provide strong support for this view. Specifically, in the absence of a retrospective cue we observed a gradual degradation in location-specific WM representations with time (Fig. 4A). However, a 100% valid retrospective cue presented immediately after the encoding display eliminated degradation in the representation of the cued disc and accelerated degradation in the representation of the uncued disc. Rapid degradation of the representation of the uncued disc is nominally consistent with studies suggesting that retrospective cues engage mechanisms that facilitate the removal of uncued items from WM (Astle et al., 2012; Kuo et al., 2012; Williams and Woodman, 2012; Williams et al., 2013). However, it is unclear whether rapid degradation reflects the operation of an active mechanism that purges irrelevant information from WM or the passive (but rapid) decay of information following the withdrawal of attention. This question awaits further scrutiny.

During VL trials, we observed a partial recovery of location information during the second half of the delay period (Fig. 6). This result dovetails with several recent empirical studies (Lewis-Peacock et al., 2012; LaRocque et al., 2013; Rose et al., 2016; Sprague et al., 2016; Wolff et al., 2017) documenting a recovery or “resurrection” of decodable stimulus information following a retrospective cue or neurostimulation. Many of these studies

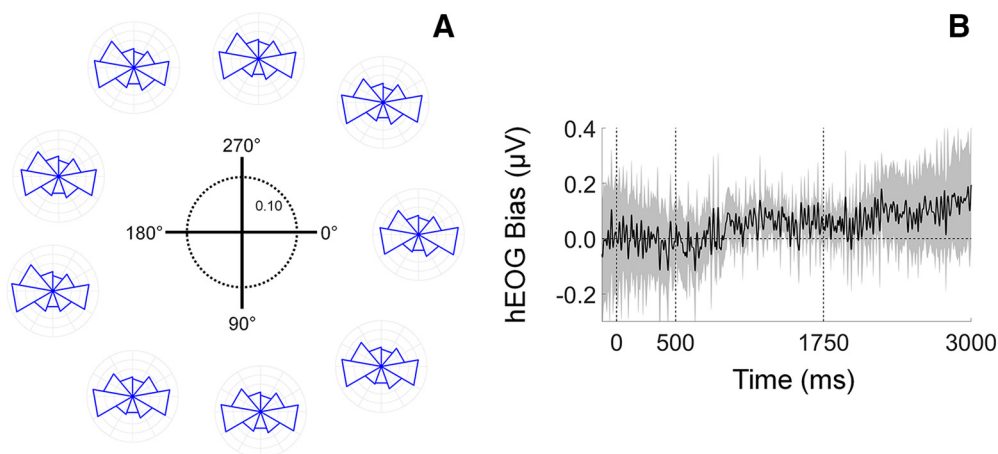


Figure 8. Location-specific reconstructions cannot be explained by subtle biases in eye position. Our analyses focused on VE trials as this is where systematic biases in eye position should be most apparent (i.e., because participants were required to remember only one location). **A**, Circular histograms showing angular estimates of eye position are plotted as a function of stimulus location. Each histogram corresponds to one of the nine possible stimulus locations (Fig. 3A). Data are scaled according to the schematic in the center of the plot. **B**, Regression coefficients describing the relationship between horizontal EOG voltage (hEOG; in microvolts) and remembered stimulus locations. Vertical dashed lines at 0 and 500 ms mark the onset and offset of the sample display, while the dashed line at 1750 ms represents the mid-point of the delay period. Shaded regions are 95% within-participant confidence intervals.

have interpreted this recovery as evidence for the existence of an additional “latent” WM system based on short-term synaptic plasticity that can be used to supplement active storage mechanisms such as sustained spiking activity (Barak and Tsodyks, 2014; Stokes, 2015). Schneegans and Bays (2017) recently challenged this conclusion by demonstrating that a neural process model based on sustained spiking activity can yield a recovery of location information in a retrospectively cued spatial WM task similar to the one used in this study (see also Sprague et al., 2016). In this model, a colored retrospective cue provides a uniform or homogeneous boost in the activity of one of two neural populations with joint selectivity for a specific color (e.g., red) and location. This boost, coupled with inhibitory interactions between the two neural populations, provides a plausible explanation for the partial recovery of location information following a retrospective cue that was observed during VL trials. However, this model would also predict a boost in location information above and beyond that seen during encoding following the presentation of the retrospective cue during VE trials (Souza et al., 2014; also see related empirical findings in the study by Rerko et al., 2014). Our findings do not support this prediction: we saw no evidence for an increase in location information following the presentation of the retrospective cue during VE trials (Fig. 5). This implies that there is an upper limit on the information content of reconstructed spatial WM representations that is determined during encoding. Further studies will be needed to explore this possibility in detail.

The absence of a boost in reconstructed spatial WM representations during VE trials (Fig. 5) conflicts with psychophysical studies, suggesting that retrospectively cued shifts of attention can strengthen WM representations over and above their original encoding strength (Rerko et al., 2014; Souza et al., 2016). It is also nominally inconsistent with findings reported by Sprague et al.

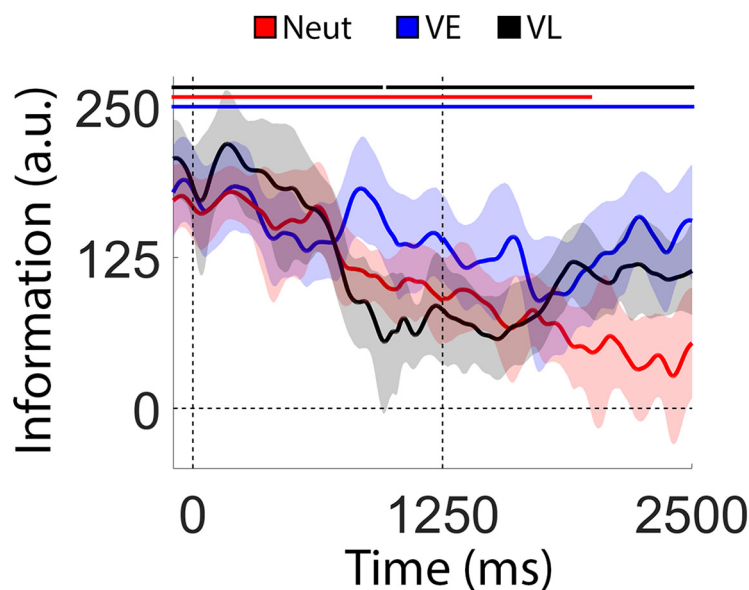


Figure 9. Synopsis of key findings. Red, blue, and black lines are reproduced from Figures 5A, 6, and 7A, respectively.

(2016), who observed an increase in the amplitude of a reconstructed spatial WM representation reconstructed from hemodynamic activity when a valid cue was presented immediately after an encoding display (compared with a neutral cue condition). However, the sluggish nature of the human hemodynamic response makes it difficult to infer the sources of this effect. Indeed, the data reported by Sprague et al. (2016) were acquired with a temporal resolution of 2250 ms, or nearly the length of the entire delay period in the current study. Thus, higher-amplitude representations during valid trials could in principle reflect a boost in reconstructed representations over and above their original encoding strength (Rerko et al., 2014; Souza et al., 2016; Schneegans and Bays, 2017) or later degradation in reconstructed representations during neutral trials. Our approach allowed us to disambiguate these possibilities by tracking changes in location-specific reconstructions with a temporal resolution on the order of tens of milliseconds.

Our findings are consistent with studies documenting links between alpha-band topography and spatial attention both

across and within visual hemifields (Rihs et al., 2007; Bahramisharif et al., 2010), as well as more recent work demonstrating that momentary changes in alpha-band topography can be used to track the locus of spatial attention with high temporal resolution (Foster et al., 2017). In addition, we have shown that alpha-band topography can be used to visualize and track changes in reconstructed WM representations following a retrospective cue. In the absence of a cue, we observed a monotonic decrease in memory strength with time. A cue presented immediately after the termination of the encoding display eliminated this decrease and a cue presented midway through the subsequent delay period partially reversed it. Collectively, our findings provide new and compelling evidence that, depending on circumstances, retrospectively cued shifts of attention can (1) prevent subsequent information loss during WM storage, (2) partially reverse prior information loss, and (3) possibly facilitate the removal of irrelevant items from WM.

References

- Astle DE, Summerfield J, Griffin I, Nobre AC (2012) Orienting attention to locations in mental representations. *Atten Percept Psychophys* 74:146–162. [CrossRef Medline](#)
- Bae GY, Luck SJ (2018) Dissociable decoding of spatial attention and working memory from EEG oscillations and sustained potentials. *J Neurosci* 38:409–422. [CrossRef Medline](#)
- Bahramisharif A, van Gerven M, Heskes T, Jensen O (2010) Covert attention allows for continuous control of brain-computer interfaces. *Eur J Neurosci* 31:1501–1508. [CrossRef Medline](#)
- Barak O, Tsodyks M (2014) Working models of working memory. *Curr Opin Neurobiol* 25:20–24. [CrossRef Medline](#)
- Bays PM, Taylor R (2018) A neural model of retrospective attention in visual working memory. *Cogn Psychol* 100:43–52. [CrossRef Medline](#)
- Blankertz B, Lemm S, Treder M, Haufe S, Müller KR (2011) Single-trial analysis and classification of ERP components—a tutorial. *Neuroimage* 56:814–825. [CrossRef Medline](#)
- Cousineau D (2005) Confidence intervals in within-subject designs: a simpler solution to Loftus and Masson's method. *Tutor Quant Methods Psychol* 1:42–45. [CrossRef](#)
- Cowan N (2001) The magical number 4 in short-term memory: a reconsideration of mental storage capacity. *Behav Brain Sci* 24:87–185. [Medline](#)
- Ester EF, Anderson DE, Serences JT, Awh E (2013) A neural measure of precision in visual working memory. *J Cogn Neurosci* 25:754–761. [CrossRef Medline](#)
- Ester EF, Sprague TC, Serences JT (2015) Parietal and frontal cortex encode stimulus-specific mnemonic representations during visual working memory. *Neuron* 87:893–905. [CrossRef Medline](#)
- Ester EF, Sutterer DW, Serences JT, Awh E (2016) Feature-selective attentional modulations in human frontoparietal cortex. *J Neurosci* 36:8188–8199. [CrossRef Medline](#)
- Foster JJ, Sutterer DW, Serences JT, Vogel EK, Awh E (2016) The topography of alpha-band activity tracks the content of spatial working memory. *J Neurophysiol* 115:168–177. [CrossRef Medline](#)
- Foster JJ, Sutterer DW, Serences JT, Vogel EK, Awh E (2017) Alpha-band oscillations enable spatially and temporally resolved tracking of cover spatial attention. *Psychol Sci* 28:929–941. [CrossRef Medline](#)
- Griffin IC, Nobre AC (2003) Orienting attention to locations in internal representations. *J Cogn Neurosci* 15:1176–1194. [CrossRef Medline](#)
- Kleiner M, Brainard D, Pelli D (2007) What's new in Psychtoolbox-3? *Perception* 36:1–16. [CrossRef](#)
- Kok P, Mostert P, de Lange FP (2017) Prior expectations induce prestimulus sensory templates. *Proc Natl Acad Sci USA* 114:10473–10478. [CrossRef Medline](#)
- Kuo BC, Stokes MG, Nobre AC (2012) Attention modulates maintenance of representations in visual short-term memory. *J Cogn Neurosci* 24:51–60. [CrossRef Medline](#)
- Landman R, Spekreijse H, Lamme VA (2003) Large capacity storage of integrated objects before change blindness. *Vision Res* 43:149–164. [CrossRef Medline](#)
- LaRocque JJ, Lewis-Peacock JA, Drysdale AT, Oberauer K, Postle BR (2013) Decoding attended information in short-term memory: an EEG study. *J Cogn Neurosci* 25:127–142. [CrossRef Medline](#)
- Lewis-Peacock JA, Drysdale AT, Oberauer K, Postle BR (2012) Neural evidence for a distinction between short-term memory and the focus of attention. *J Cogn Neurosci* 24:61–79. [CrossRef Medline](#)
- Lins OG, Picton TW, Berg P, Scherg M (1993) Ocular artifacts in EEG and event-related potentials I: scalp topography. *Brain Topogr* 6:51–63. [CrossRef Medline](#)
- Loftus GR, Masson MEJ (1994) Using confidence intervals in within-subject designs. *Psych Bull Rev* 1:476–490. [CrossRef Medline](#)
- Luck SJ, Vogel EK (2013) Visual working memory capacity: from psychophysics and neurobiology to individual differences. *Trends Cogn Sci* 17:391–400. [CrossRef Medline](#)
- Ma WJ, Husain M, Bays PM (2014) Changing concepts of working memory. *Nat Neurosci* 17:347–356. [CrossRef Medline](#)
- Makovsik T, Jiang YV (2007) Distributing versus focusing attention in visual short-term memory. *Psychon Bull Rev* 14:1072–1078. [CrossRef Medline](#)
- Makovski T, Watson LM, Koutstall W, Jiang YV (2010) Method matters: systematic effects of testing procedure on visual working memory. *J Exp Psychol Learn Mem Cogn* 36:1466–1479. [CrossRef Medline](#)
- Maris E, Oostenveld R (2007) Nonparametric statistical testing of EEG and MEG data. *J Neurosci Methods* 164:177–190. [CrossRef Medline](#)
- Matsukura M, Luck SJ, Vecera SP (2007) Attention effects during visual short-term memory maintenance: protection or prioritization? *Percept Psychophys* 69:1422–1434. [CrossRef Medline](#)
- Myers NE, Walther L, Wallis G, Stokes MG, Nobre AC (2015) Temporal dynamics of attention during encoding versus maintenance of working memory: complementary views from event-related potentials and alpha-band oscillations. *J Cogn Neurosci* 27:492–508. [CrossRef Medline](#)
- Myers NE, Stokes MG, Nobre AC (2017) Prioritizing information during working memory: beyond sustained internal attention. *Trends Cogn Sci* 21:449–461. [CrossRef Medline](#)
- Pertsov Y, Bays PM, Joseph S, Husain M (2013) Rapid forgetting prevented by retrospective attention. *J Exp Psychol Hum Percept Perform* 39:1224–1231. [CrossRef Medline](#)
- Poch C, Campo P, Barnes GR (2014) Modulation of alpha and gamma oscillations related to retrospectively orienting attention within working memory. *Eur J Neurosci* 40:2399–2405. [CrossRef Medline](#)
- Rerko L, Souza AS, Oberauer K (2014) Retro-cue benefits in working memory without sustained focal attention. *Mem Cognit* 42:712–728. [CrossRef Medline](#)
- Rihs TA, Michel CM, Thut G (2007) Mechanisms of selective inhibition in visual spatial attention are indexed by alpha-band EEG synchronization. *Eur J Neurosci* 25:603–610. [CrossRef Medline](#)
- Rose NS, LaRocque JJ, Riggall AC, Gosseries O, Starrett MJ, Meyerling EE, Postle BR (2016) Reactivation of latent working memories with transcranial magnetic stimulation. *Science* 354:1136–1139. [CrossRef Medline](#)
- Samaha J, Sprague TC, Postle BR (2016) Decoding and reconstructing the focus of spatial attention from the topography of alpha-band oscillations. *J Cogn Neurosci* 28:1090–1097. [CrossRef Medline](#)
- Schneegans S, Bays PM (2017) Restoration of fMRI decodability does not imply latent working memory states. *J Cogn Neurosci* 29:1977–1994. [CrossRef Medline](#)
- Shannon CE (1948) A mathematical theory of communication. *Bell Syst Tech J* 27:379–423. [CrossRef](#)
- Souza AS, Oberauer K (2016) In search of the focus of attention in working memory: 13 years of the retro-cue effect. *Atten Percept Psychophys* 78:1839–1860. [CrossRef Medline](#)
- Souza AS, Rerko L, Oberauer K (2014) Unloading and reloading working memory: attending to one item frees capacity. *J Exp Psychol Hum Percept Perform* 40:1237–1256. [CrossRef Medline](#)
- Souza AS, Rerko L, Oberauer K (2016) Getting more from visual working memory: retro-cues enhance retrieval and protect from visual interference. *J Exp Psychol Hum* 42:890–910. [CrossRef Medline](#)
- Sprague TC, Ester EF, Serences JT (2014) Reconstructions of information in visual spatial working memory degrade with memory load. *Curr Biol* 24:2174–2180. [CrossRef Medline](#)
- Sprague TC, Saproo S, Serences JT (2015) Visual attention mitigates information loss in small- and large-scale neural codes. *Trends Cogn Sci* 19:215–226. [CrossRef Medline](#)

- Sprague TC, Ester EF, Serences JT (2016) Restoring latent visual working memory representations in human cortex. *Neuron* 91:694–707. [CrossRef Medline](#)
- Stokes MG (2015) “Activity-silent” working memory in prefrontal cortex: a dynamic coding perspective. *Trends Cogn Sci* 19:394–405. [CrossRef Medline](#)
- Sutterer DW, Foster JJ, Serences JT, Vogel EK, Awh E (2018) Alpha-band oscillations track the retrieval of precise spatial representations from long-term memory. *bioRxiv* 207860. Advance online publication. Retrieved February 18, 2018. [CrossRef](#)
- van Ede F (2017) Mnemonic and attentional roles for states of attenuated alpha oscillations in perceptual working memory: a review. *Eur J Neurosci*. Advance online publication. Retrieved August 20, 2018. doi:10.1111/ejn.13759.
- van Moorselaar D, Foster JJ, Sutterer DW, Theeuwes J, Olivers CNL, Awh E (2018) Spatially selective alpha oscillations reveal moment-by-moment trade-offs between working memory and attention. *J Cogn Neurosci* 30:256–266. [CrossRef Medline](#)
- Vogel EK, McCollough AW, Machizawa MG (2005) Neural measures reveal individual differences in controlling access to working memory. *Nature* 438:500–503. [CrossRef Medline](#)
- Williams M, Woodman GF (2012) Directed forgetting and directed remembering in visual working memory. *J Exp Psychol Learn Mem Cogn* 38:1206–1220. [CrossRef Medline](#)
- Williams M, Hong SW, Kang MS, Carlisle NB, Woodman GF (2013) The benefit of forgetting. *Psychon Bull Rev* 20:348–355. [CrossRef Medline](#)
- Wolff MJ, Jochim J, Akyürek EG, Stokes MG (2017) Dynamic hidden states underlying working-memory-guided behavior. *Nat Neurosci* 20:864–871. [CrossRef Medline](#)