Information recovery following a retrospective cue decreases with time.

Asal Nouri[1], Edward F. Ester[1-3*]

[1]FAU Brain Institute, Florida Atlantic University
[2]Department of Psychology, Florida Atlantic University
[3]Center for Complex Systems & Brain Sciences, Florida Atlantic University

[*]correspondence:

Edward Ester
Department of Psychology, Florida Atlantic University
777 Glades Rd.
BS-12, RM-101
Boca Raton, FL., 33431, USA
eester@fau.edu

Author Contributions: A.N. and E.F.E. conceived and designed the experiment. A.N. collected data; A.N. and E.F.E. analyzed data. A.N. and E.F.E. wrote the paper.

Abstract Word Count: 213/250

Total Word Count (inc. references): 4444

Total Word Count (exc. References): 3750

Figure/Table Count: 5/0

Data & Software available on the Open Sciences Framework: https://osf.io/vqc26/

Abstract

Working memory (WM) performance can be enhanced by an informative cue presented during storage. This effect, termed a retrocue benefit, can be used to explore how observers prioritize information stored in WM to guide behavior. Recent studies have demonstrated that neural representations of task-relevant memoranda are strengthened following the appearance of a retrocue, suggesting that participants can consult alternative information stores to supplement active memory traces. Here, we sought to better understand the nature of these memory store(s) by asking whether they are subject to the same temporal degradation seen in active memory representations. We reconstructed and quantified representations of remembered positions from alpha-band EEG activity recorded while participants performed a retrospectively cued spatial WM task, and varied the temporal interval separating the encoding display and retrocue. Although we observed a partial recovery of location information in all cue conditions, the magnitude of recovery was linearly and inversely related to the timing of the retrocue. This suggests that participants' ability to supplement active memory representations with information from additional memory stores is not static: the information maintained in these stores may be subject to temporal degredation, or the stores themselves may become more difficult to access with time.

Acknowledgements: Funding for this project was provided by a State University System of Florida startup award to E.F.E. The authors wish to thank Brian Escobar for assistance with data collection.

Working memory (WM) enables the representation of information no longer present in the sensorium. This system is a core component of general cognitive function, as evidenced by robust correlations between measures of WM ability and aptitudes such as IQ, scholastic achievement, and reading ability (e.g., Daneman & Carpenter, 1980; Engle et al., 1999). Despite this importance, WM is famously burdened with a limited capacity (Luck & Vogel, 2013; Ma et al., 2014). Thus, mechanisms of selective attention are needed to gate access to WM and to prioritize existing WM representations for behavioral output.

Attentional prioritization in WM can be studied using retrospective cues (see Souza & Oberauer, 2016; Myers et al. 2017; for recent comprehensive reviews). In a typical retrospective cue experiment, participants are asked to store an array of objects (e.g., colored squares) for subsequent report. At some point during the ensuing delay period, a cue indicates which of the original memory items is most likely to be tested. Relative to a neutral-cue or no-cue baseline, valid cues yield superior memory performance. This effect is commonly referred to as a retrocue benefit.

Several mechanisms may be responsible for retrocue benefits, including (but not limited to): (a) an attentional enhancement of the cued memory representation (Rerko, Souza, & Oberauer, 2014; Souza, Rerko, & Oberauer, 2015), (b) insulation of the cued memory representation from subsequent decay or interference (Pertzov, Bays, Joseph, & Husain, 2013), (c) removal of non-cued representations from memory (Souza & Oberauer, 2016), or (d) some combination of the above. A recent study (Ester, Nouri, Rodriguez, 2018) sought to evaluate these alternatives by recovering and tracking neural representations of memoranda from human EEG activity. Participants in this

study performed a retrospectively cued spatial WM task that required them to remember the spatial positions of two discs over a blank delay. During valid trials, a 100% reliable retrocue informed participants which disc's location would be probed at the end of the trial. Importantly, valid cues were presented either immediately after the offset of the encoding display (valid-early trials) or midway through the subsequent blank delay (valid-late trials).

To examine the effects of retro-cues on spatial WM representations, Ester et al. (2018) computed an estimate of location-specific information for the cued and non-cued disc locations by applying an inverted encoding model to spatiotemporal patterns of induced alpha-band activity over occipitoparietal electrode sites during the blank delay period (e.g., Foster et al. J Neurophysiol 2016). During neutral trials, where no cue was presented, the authors observed a monotonic decrease in location-specific information over the course of the delay period. However, this decrease was eliminated during valid early trials and partially reversed during valid late trials. The results of this study suggest that retrocue benefits may reflect the operation of several different mechanisms, including the removal of non-cued representations from memory and the insulation of cued representations from temporal decay or interference.

In keeping with earlier studies (e.g., Sprague et al., 2016; Wolff et al., 2017), Ester et al. (2018) interpreted the partial recovery of location-specific information observed during valid late trials as evidence that participants were able to supplement active memory representations indexed by alpha-band activity with other information sources not indexed by this activity, including but not limited to "activity silent" WM or long-term memory. In this study, we asked whether participants' ability to exploit these

additional memory stores is time-dependent. Specifically, we wondered whether information stored in “offline” memory systems degrades or becomes more difficult to access over time, much in the same way that active memory representations indexed by alpha-band activity degrade during neutral trials. If so, then one would expect that the extent of retrocue-driven recovery in estimates of location information would decrease as the temporal interval separating the encoding display and retrocue increases. We tested this possibility by replicating the neutral and VL conditions from Ester et al. (2018) while varying the timing of the retrocue between 1.0, 1.5, and 2.0 sec after the encoding display As expected, we found that degree of information recovery was linearly and inversely related to the interval separating the retrocue and the encoding display, suggesting that representations stored in “offline” memory systems degrade or become more difficult to access over time.

Methods

Participants. Thirty-six (36) volunteers were recruited from the Florida Atlantic University undergraduate community and tested in a single 2.5-hour session. All experimental procedures were approved by the local institutional review board. Participants gave both written and oral informed consent, self-reported normal or corrected-to-normal visual acuity, and were compensated with course credit or monetary remuneration ($15/h in Amazon gift cards). Data from two participants were discarded due to large numbers of electrooculogram artifacts (over 50% of trials). The data reported here reflect the remaining 34 participants.

Testing Environment. Participants were tested in a dimly lit, sound-attenuated recording chamber. Stimulus displays were generated in MATLAB and rendered on a 17-inch dell CRT monitor cycling at 85 Hz (1024 x 768 pixel resolution) via Psychophysics Toolbox software extensions (Kleiner et al., 2007). Participants were seated approximately 60 cm from the display (head position was not constrained) and made responses using a standard computer keyboard and mouse.

Behavioral Task. 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 8 equally spaced positions (22.5° to 337.5° in 45° increments) along the perimeter of an imaginary circle (radius 6° visual angle) centered at the fixation point (Figure 1B). A small amount of jitter (±10° 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 0.5 sec, at which point the fixation point changed colors from black to purple. The fixation point either remained purple for the ensuing 3.0 sec delay period (neutral trials) or changed colors from purple to either blue or red 1.0, 1.5, or 2.0 sec after termination of the sample display (valid trials). When presented, valid cues indicated which disc participant would be probed to report at the end of the trial with 100% reliability. 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 3.0 sec 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 was given at the end of each block. Participants completed 6 (N = 1), 11 (N = 1) or 12 (N = 32) blocks of 56 trials.

EEG Recording and Preprocessing. Continuous EEG was recorded from 63 Ag/Ag-Cl^- scalp electrodes via a Brain Products actiCHamp amplifier. An additional electrode was placed over the right mastoid. Data were recorded with a right mastoid reference and later re-referenced to the algebraic mean of the left and right mastoids (10-20 site TP9 served as the left mastoid reference). The horizontal and vertical electrooculogram (EOG) was recorded from electrodes placed on the left and right canthi and above and below the right eye, respectively. All electrode impedances were kept below 15 kΩ, and recordings were digitized at 1000 Hz. Recorded data were bandpass filtered from 1 to 25 Hz (3rd order zero-phase forward and reverse Butterworth filters), epoched from a period spanning -500 to +4000 ms relative to the start of each trial, and baseline

corrected from -125 to 0 ms. Trials contaminated by blinks or horizontal eye movements greater than ~2.5° (assuming a normative voltage threshold of 16 μV/°; 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 S.E.M.) of 17.69% (±2.07%) trials and 1.11 (±0.27) electrodes.

Inverted Encoding Model. Following earlier work (e.g., Foster et al. 2016; Ester et al. 2018) we used an inverted encoding model 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 17 occipitoparietal electrode sites: O1, O2, Oz, PO7, PO3, POz, PO4, PO8, P7, P5, P3, P1, Pz, P2, P4, P6, and P8. To isolate alpha-band activity, the raw EEG time series at each electrode was bandpass filtered from 8-12 Hz (zero-phase forward and reverse filters as implemented by EEGLAB's "eegfilt" function), yielding a real-valued signal $f(t)$. The analytic representation of $f(t)$ was obtained by applying a Hilbert transformation:

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

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 8 location-selective channels, each with an idealized tuning curve (a half-wave rectified cosine raised to the 8th power). The maximum response of each channel was

normalized to 1, thus 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 trials (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 form:

$$B = WC + N$$

where B is a *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" data set containing an equal number of trials from each stimulus location (i.e., 45-360° in 45° steps) condition. We first identified the location φ with the fewest *r* repetitions in the full data set after EOG artifact removal. Next, we constructed a training data set 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 eight possible stimulus locations (ignoring cue condition; i.e., the training data set contained a mixture of neutral and valid trials). The training data set was used to compute a weight for each channel C_i via least-squares estimation:

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

where $C_{trn,i}$ 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):

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

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

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

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

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

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

Next, we constructed a “test” data set B_{tst} (m electrodes by n trials) containing data from all trials not included in the training data set and estimated trial-by-trial

channel responses C_{tst} (k channels x n trials) from the filter matrix V and the test data set:

$$C_{tst} = V^T B_{tst}$$

Trial-by-trial channel responses were interpolated to 360°, circularly shifted to a common center (0°, by convention), and sorted by cue condition (e.g., neutral vs. valid). To quantify the effects of retrospective cues reconstructed representations, we obtained an estimate of representation strength by converting the (averaged) channel response estimates for each cue condition to polar form and projected them onto a vector with angle 0°:

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

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 50 times, and unique (randomly chosen) subsets of trials were used to define the training and test data sets during each permutation. The results were then averaged across permutations.

Eye Movement Control Analyses. Although we excluded trials contaminated by horizontal eye movements, small but reliable biases in eye position towards the location(s) of the remembered disc(s) could nevertheless contribute to reconstructions

of stimulus location. We tested this possibility by regressing trial-by-trial horizontal EOG recordings (in μV) onto the horizontal position of the cued disc (see Foster et al., 2016). We restricted our analysis to valid trials on the assumption that eye movements would be most prevalent when participants were explicitly instructed which of the two discs to remember. Positive regression coefficients 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.

Statistical Comparisons. All statistical comparisons were based on non-parametric permutation (bootstrap) tests. For example, to compare the magnitude of the retrocue benefit during the 1.0 and 1.5 cue delay conditions, we randomly selected (with replacement) and averaged data from 34 of 34 participants in each condition, then computed the difference between the two means. This procedure was repeated 10,000 times, yielding a 10,000-element distribution of average differences. Finally, we computed the number of permutations where the cue benefit was opposite the expected direction, yielding an empirical p-value. Where appropriate, p-values were false-discovery-rate (FDR) corrected for multiple comparisons using the Benjamini & Hochstein (1995) procedure.

Within-participant Variability. We report estimates of within-participant variability (e.g., 95% within-participant confidence intervals) throughout the paper. These estimates discard subject variance (e.g., overall differences in response strength) and instead reflect variance related to the subject by condition(s) interaction term(s) (e.g., variability in response strength across experimental conditions; Loftus & Masson, 1994; Cousineau 2005). We used the approach described by Cousineau (2005): raw data

(e.g., location information or channel response estimates) were de-meaned 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

Behavioral Performance. Average absolute report error estimates are plotted as a function of cue condition in Figure 1D. Report errors were reliably lower during all three cue conditions relative to neutral trials (false-discovery-rate-corrected bootstrap tests, $p < 0.006$). Direct comparisons across the valid cue conditions revealed a reliable difference between the 1.0 and 2.0 s conditions (FDR-corrected bootstrap test; $p = 0.005$). The differences between 1.0 and 1.5 sec conditions and the 1.5 and 2.0 sec conditions were not significant ($p > 0.125$). There was a significant linear increase in report error estimates as a function of cue delay (M = 9.85°, 10.21°, and 10.56° for the 1.0, 1.5, and 2.0 sec due conditions, respectively; slope = +0.70°/sec, $p = 0.001$). Thus, participants' report errors were partially contingent on the timing of the retrocue, with

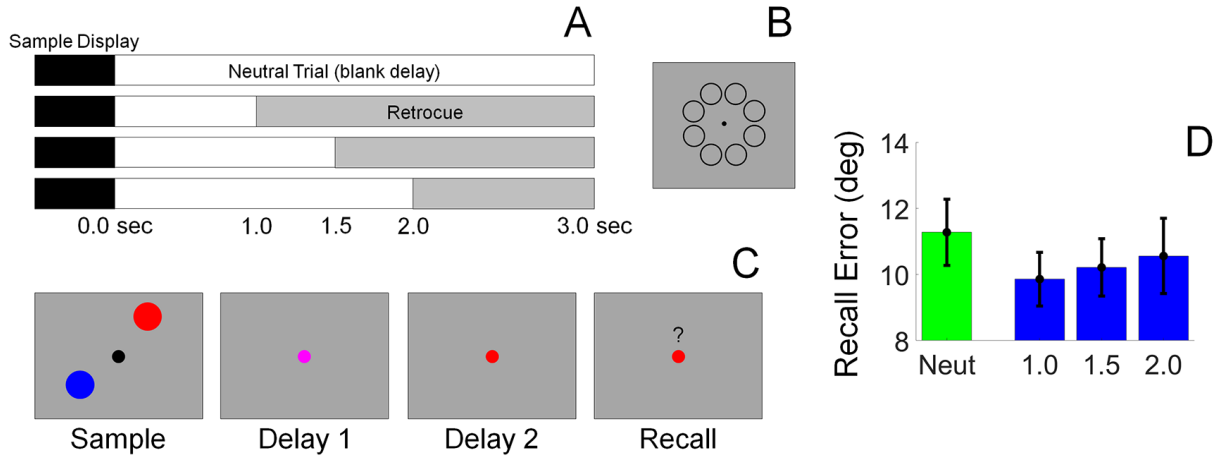

Figure 1. Task Design and Behavioral Data. (A) Schematic showing the different cue conditions. On 50% of trials, the sample display was followed by a 3.0 sec blank delay (neutral trials). On the remaining 50% of trials, a 100% valid retrocue was rendered 1.0, 1.5, or 2.0 sec after termination of the sample display. In all cases, the total delay period length was fixed at 3.0 seconds. (B) Schematic showing possible disc positions. (C) Schematic showing the sequence of events during each trial (Note: stimuli are not drawn to scale; see Methods for information regarding display and stimulus geometry). (D) Behavioral performance, quantified as average absolute recall error (i.e., the average absolute difference between the polar location reported by the participant and the polar location of the probed disc)

superior performance for earlier relative to later retrocues (Pertzov et al., 2013; Souza et al. 2016).

Reconstructing location-specific WM representations from alpha-band activity. Following earlier studies (Foster et al., 2016; Ester et al., 2018) we used a spatial inverted encoding model to reconstruct the locations of the cued and uncued discs from spatiotemporal patterns of posterior alpha band power. Specifically, we modeled induced alpha power recorded at each electrode site as a weighted sum of eight location-selective channels, each with an idealized response function. The resulting weights were used to calculate a predicted response for each channel given patterns of induced alpha power recorded during an independent data set. Trial-by-trial reconstructions were circularly shifted to a common center (0° by convention), yielding a single time-resolved reconstruction for each participant. We quantified these reconstructions by converting them to polar form and projecting them onto a unit vector with angle 0°. The resultant vector length was interpreted as a measure of total location information.

Estimates of location information are plotted as a function of cue condition in Figure 2. In the absence of a retrospective cue (neutral trials), estimates of representation strength declined monotonically throughout the delay period (-21.04 units/sec; $p < 0.0001$). However, visual inspection of Figure 2 suggests that this decline was partially reversed following the appearance of a retrocue. To formally evaluate this possibility, we compared average location information estimates for each valid cue condition over an interval spanning 0-500 ms after cue onset (i.e., 1.0 to 1.5 sec for the

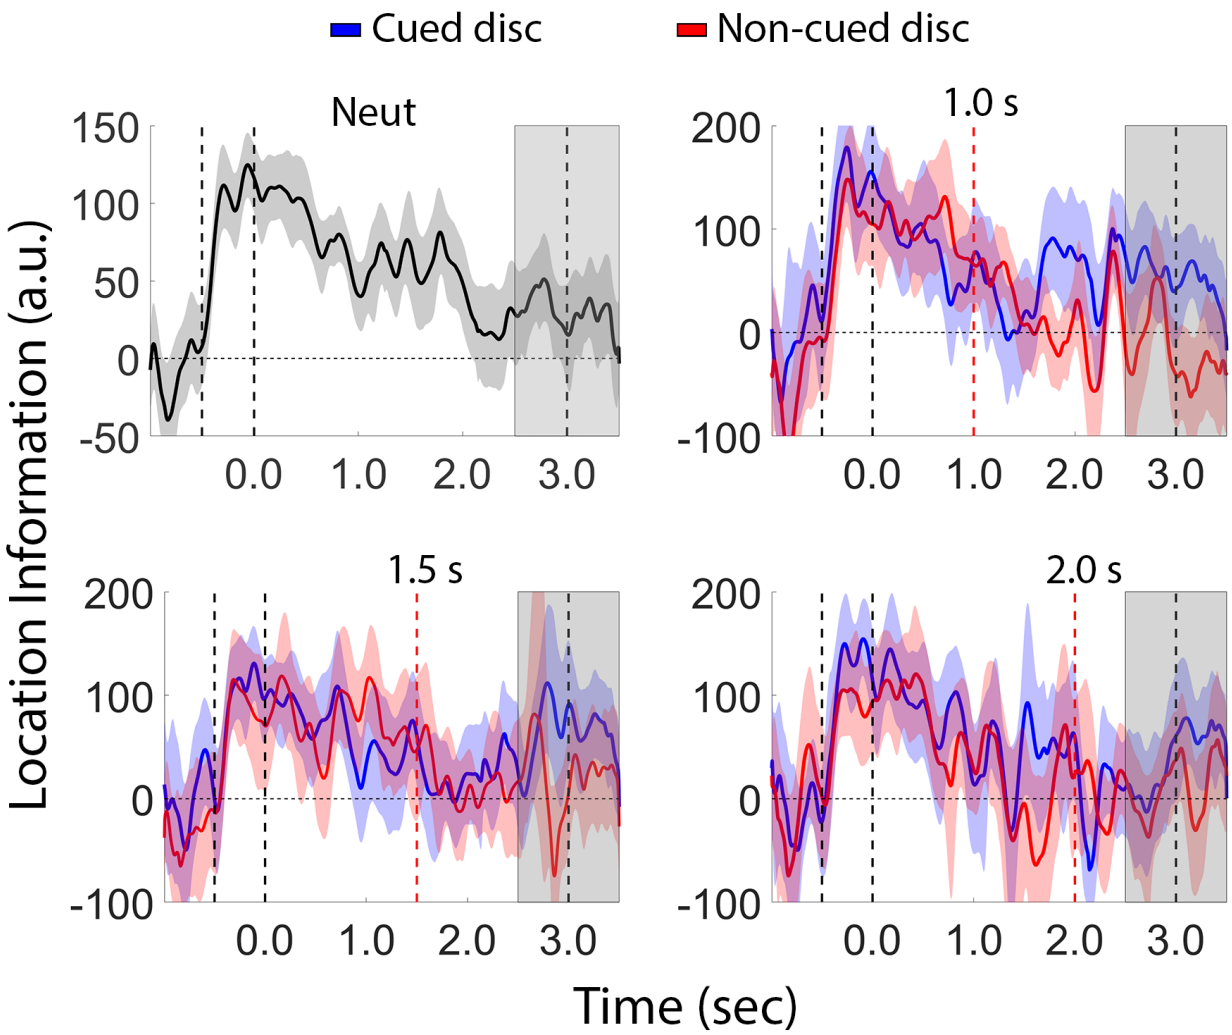

Figure 2. Estimates of Location Information Recovered from EEG Activity. Black vertical dashed bars at 0.0 and 3.0 sec indicate the start and end of the blank delay period. Red vertical dashed bars indicate the onset of the retrocue during valid trials. The shaded region from 2.5 to 3.5 seconds was used in later statistical comparisons (see Figures 3 and 4). The shaded area around each plot depicts the 95% within-participant confidence interval of the mean.

1.0 sec cue condition, 1.5 to 2.0 sec for the 1.5 sec cue condition, and 2.0 to 2.5 sec for the 2.0 sec cue condition) with average information estimates measured over a fixed interval spanning 2.5 to 3.5 sec after termination of the sample display[1]. The resulting averages are plotted in Figure 3. There was a significant recovery of location information for the cued disc during the 1.5 and 2.0 sec cue conditions (FDR-corrected

[1] The 0-500 ms window following cue onset was chosen based on behavioral studies suggesting that it takes participants about 300-500 ms to process and utilize a retrospective cue (e.g., Pertzov et al., 2013).

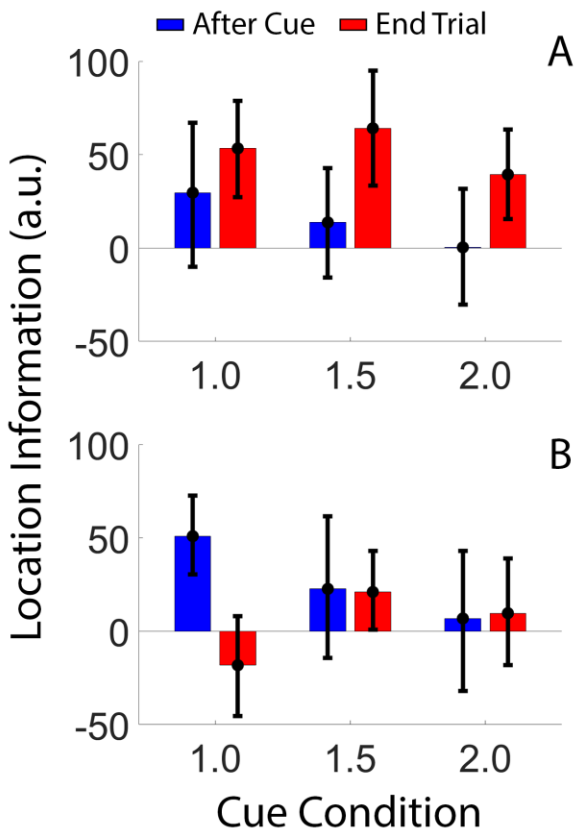

Figure 3. Recovery of Information Following a Retrocue. Time-based changes in location information for the cued disc (A) and the non-cued disc (B). A significant increase in information was observed for the cued disc during the 1.5 and 2.0 sec cue conditions. Error bars depict the 95% within-participant confidence interval of the mean.

p-values < 0.05), but not during the 1.0 sec cue condition (p = 0.18). There was a significant decrease in location information for the uncued disc during the 1.0 sec cue condition (p < 0.0001) but not during the 1.5 and 2.0 sec cue conditions (p > 0.51; Figure 3B). Thus, the current findings are consistent with those reported by Ester et al. (2018): memory representations reconstructed from alpha-band EEG activity degrade with time, but a valid retrocue can partially reverse earlier degradation.

The recovery of information following a retrospective cue declines with time. Recovery of information following a retrocue suggests that participants can supplement active memory representations with information stored in memory systems not indexed by alpha-band activity, such as activity-silent short-term memory or long-term memory. The goal of this study was to examine whether information held in these alternative stores degrades or becomes more difficult to access over time. If so, then the amount of information recovery seen in representations of memoranda reconstructed from alpha band activity (Figures 2 and 3) should decrease as the temporal interval separating the encoding display from the retrocue increases. We tested this possibility by calculating

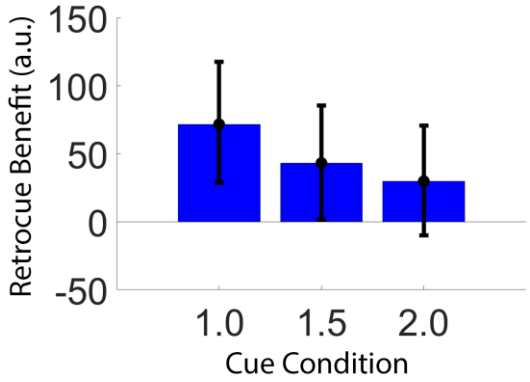

Figure 4. Retrocue Benefits Decrease with Time. Error bars depict the 95% within-participant confidence interval of the mean.

the total retrocue benefit – defined as the difference in location information estimates for the cued and non-cued discs over an interval spanning 2.5 to 3.5 sec following termination of the sample display – as a function of cue condition. The resulting estimates are plotted in Figure 4. We observed a monotonic decrease in the magnitude of the retrocue benefit with increasing sample-to-cue inter-stimulus-intervals (M = 71.58, 43.09, and 29.71 units for the 1.0, 1.5, and 2.0 sec cue conditions; slope = -41.53 units/sec; p = 0.043), suggesting that participants ability to exploit information stored in other memory stores is also time-limited.

Ruling out contributions from eye movements. We identified and excluded trials contaminated by horizontal eye movements greater than 2.5° (based on a normative threshold of 16 µV/°; Lins et al. 1993). However, smaller and consistent eye movement patterns across trials could also contribute to the location reconstructions reported here. We examined this possibility in greater detail by regressing time-resolved estimates of horizontal EOG activity onto remembered

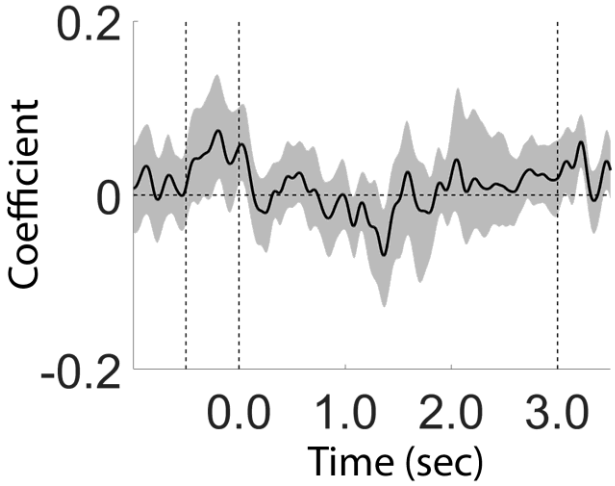

Figure 5. Effect of Eye Position on Stimulus Reconstructions. Black vertical dashed lines at 0.0 and 3.0 depict the start and end of the delay period. Shaded regions depict the 95% within-participant confidence interval of the mean.

stimulus locations. As shown in Figure 5, the regression coefficients linking eye position with remembered locations were indistinguishable from 0 for the duration of each trial, suggesting that eye movements were not a major determinant of location reconstructions.

Discussion

Working memory performance can be improved by an informative retrocue presented during storage. This effect, termed a retrocue benefit, can be used to study mechanisms of attentional priority in WM. Recent neuroimaging (EEG and fMRI) studies have identified a partial recovery in neural representations of cued memory items following a retrospective cue (Sprague et al., 2016; Rose et al., 2016; Wolff et al., 2017; Christophel et al., 2018; Ester et al., 2018), suggesting that participants can query information held in “activity-silent” memory stores to supplement active memory representations. Here, we asked whether these additional memory stores are subject to the same degradation (or become more difficult to access) over time in the same way as active memory representations (e.g., Figure 2). To test this possibility, we used an inverted encoding model to recover and quantify the fidelity of location-specific representations from EEG activity recorded while participants performed a retrospectively cued spatial WM task. In addition, we systematically varied the timing of the retrocue relative to the termination of the encoding display from between 1.0 to 2.0 seconds. Consistent with earlier findings (e.g., Ester et al., 2018) we observed a recovery in location information estimates following the appearance of a retrospective cue (Figures 2-3). However, the magnitude of this recovery was inversely and linearly related to the interval separating the termination of the encoding display and the onset of the retrospective cue (i.e., less recovery for longer intervals; Figure 4). This implies that the information stored in activity silent mechanisms is subject to temporal degradation (either due to decay or inter-item interference), or that it becomes more difficult to access over time.

Candidate mechanisms for alternative memory stores include “activity-silent” working memory (e.g., based on short-term synaptic plasticity; Mongillo, Barak, & Tsodyks, 2008; Stokes, 2015; Wolff et al., 2017) or long-term memory (Sutterer et al., 2018). On the one hand, there is ample evidence that long-term memory systems can be used to complement or supplement active memory storage when the task allows (e.g., Carlisle et al., 2011). However, behavioral studies have found little to no evidence for time-based degradation in visual long-term memory. For example, in some instances recall of visual features from long-term memory has been found to match performance during recall from WM (Brady et al., 2008; Brady et al., 2013). Based on these findings, we are skeptical that retrieval from long-term memory can account for the data reported here. On the other hand, models of activity-silent working memory information propose that task-relevant representations are maintained through short-term synaptic plasticity (Zucker & Regehr, 2002). An influential modeling study (Mongillo et al., 2008) suggests that calcium kinetics in prefrontal cortex could act as a short-term memory buffer with a time constant of about 1-2 seconds. In the absence of continuous spiking input, representations encoded in this format would be expected to degrade with time. This account naturally accommodates the time-based decreases in the magnitude of information recovery following a retrocue reported here.

References

Blankertz B, Lemm S, Treder M, Gaufe S, Müller K-R. Single-trial analysis and classification of ERP components – A tutorial. NeuroImage 56:814-825 (2011)

Brady TF, Konkle T, Alvarez GA, Oliva A. Visual long-term memory has a massive storage capacity for object details. Proc Natl Acad Sci USA 105:14325-14329 (2008)

Brady TF, Konkle T, Gill J, Oliva A, Alvarez GA. Visual long-term memory has the same limit on fidelity as visual working memory. Psychol Sci 24:981-990 (2013)

Carlisle NB, Arita JT, Pardo D, Woodman GF. Attentional templates in visual working memory. J Neurosci 31:9315-9322 (2011)

Christophel TB, Iamshchinina P, Yan C, Allefeld C, Haynes JD. Cortical specialization for attended versus unattended working memory. Nat Neurosci 21:494-496. (2018)

Cousineau D. Confidence intervals in within-subject designs: A simpler solution to Loftus and Masson's method. Tutorials in Quantitative Methods for Psychology. 1:42-45 (2005)

Daneman M, Carpenter P.A. Individual differences in working memory and reading. *Journal of Verbal Learning & Verbal Behavior* 19:450-466 (1980)

Engle RW, Tuholski SW, Laughlin JE, Conway ARA. Working memory, short-term memory, and general fluid intelligence: A latent variable approach. *J Exp Psychol Gen* 128:309-321 (1999)

Ester EF, Nouri A, Rodriguez L. Retrospective cues mitigate information loss in human cortex during working memory storage. *J Neurosci* 38:8538-8458 (2018)

Foster JJ, Sutterer DW, Serences JT, Vogel EK, Awh E. The topography of alpha-band activity tracks the content of spatial working memory. J Neurophysiol 115:168-177 (2016)

Kleiner M, Brainard D, Pelli D. What's new in Psychtoolbox-3? *Perception* (2007)

Kok P, Mostert P, de Lange FP. Prior expectations induce prestimulus sensory tempulates. *Proc Natl Acad Sci USA* (2017)

Kuo B, Stokes MG, Nobre AC. Attention modulates maintenance of representations in visual short-term memory. J Cogn Neurosci 24:51-60. (2012)

Lins OG, Picton TW, Berg P, Scherg M. Ocular artifacts in EEG and event-related potentials I: Scalp topography. *Brain Topography* 6:51-63. (1993)

Loftus GR, Masson MEJ. Using confidence intervals in within-subject designs. Psych Bull Rev 1:476-490. (1994)

Luck SJ, Vogel EK. Visual working memory capacity: Form psychophysics and neurobiology to individual differences. *Trends Cogn Sci* 17:391-400 (2013)

Ma WJ, Husain M, Bays PM. Changing concepts of working memory. *Nat Neurosci* 17:347-356. (2014)

Maris E, Oostenveld R. Nonparametric statistical testing of EEG and MEG data. J Neurosci Methods 164: 177-190 (2007)

Mongillo G, Barak O, Tsodyks M. Synaptic theory of working memory. Science 319:1543-1546 (2008)

Myers NE, Walther L, Wallis G, Stokes MG, Nobre AC. Temporal dynamics of attention during encoding versus maintenance of working memory: Complementary views

from eventrelated potentials and alpha-band oscillations. J Cogn Neurosci 27:492-508. (2015)

Myers NE, Stokes MG, Nobre AC. Prioritizing information during working memory: Beyond sustained internal attention. Trends Cogn Sci 21:449-461 (2017)

Pertzov Y, Bays PM, Joseph S, Husain M. Rapid forgetting prevented by retrospective attention. J Exp Psychol: Hum 39:1224-1231 (2013)

Rerko L, Souza AS, Oberauer K. Retro-cue benefits in working memory without sustained focal attention. Memory & Cognition 42:712-728 (2014)

Rose NS, LaRocque JJ, Riggall AC, Gosseries O, Starrett MJ, *et al*. Reactivation of latent working memories with transcranial magnetic stimulation. *Science* 354:1136-1139 (2016)

Souza AS, Rerko L, Oberauer K. Refreshing memory traces: Thinking of an item improves retrieval from visual working memory. Annals of the New York Academy of Sciences, 1339:20-31 (2015)

Souza AS, Rerko L, Oberauer K. Getting more from visual working memory: Retro-cues enhance retrieval and protect from visual interference. JEP HPP (2016)

Souza AS, Oberauer K. In search of the focus of attention in working memory. 13 years of the retro-cue effect. *Atten Percept Psychophys* 78:1839-1860 (2016)

Sprague TC, Ester EF, Serences JT. Restoring latent visual working memory representations in human cortex. Neuron 91:694-707 (2016)

Stokes MG. Activity-silent working memory in prefrontal cortex: A dynamic coding framework. Trends Cogn Sci 19:394-405 (2015)

Sutterer DW, Foster JJ, Serences JT, Vogel EK, Awh E. Alpha-band socillations track the retrieval of precise spatial representations from long-term memory. bioRxiv. doi: https://doi.org/10.1101/207860

Williams M, Woodman GF. Directed forgetting and directed remembering in visual working memory. J Exp Psychl Learn Mem Cogn 38:1206-1220 (2012)

Wolff MJ, Jochim J, Akyürek EG, Stokes MG. Dynamic hidden states underlying working memory-guided behavior. Nat Neurosci 20:864-871 (2017).

Zucker RS, Regehr WG. Short-term synaptic plasticity. Annu Rev Physiol 64:355-405 (2002).