

Rapid Forgetting Prevented by Retrospective Attention Cues

Yoni Pertzov, Paul M. Bays, and Sabine Joseph
University College London

Masud Husain
University College London and University of Oxford

Recent studies have demonstrated that memory performance can be enhanced by a cue which indicates the item most likely to be subsequently probed, even when that cue is delivered seconds after a stimulus array is extinguished. Although such retro-cuing has attracted considerable interest, the mechanisms underlying it remain unclear. Here, we tested the hypothesis that retro-cues might protect an item from degradation over time. We employed two techniques that previously have not been deployed in retro-cuing tasks. First, we used a sensitive, continuous scale for reporting the orientation of a memorized item, rather than binary measures (change or no change) typically used in previous studies. Second, to investigate the stability of memory across time, we also systematically varied the duration between the retro-cue and report. Although accuracy of reporting uncued objects rapidly declined over short intervals, retro-cued items were significantly more stable, showing negligible decline in accuracy across time and protection from forgetting. Retro-cuing an object's color was just as advantageous as spatial retro-cues. These findings demonstrate that during maintenance, even when items are no longer visible, attention resources can be selectively redeployed to protect the accuracy with which a cued item can be recalled over time, but with a corresponding cost in recall for uncued items.

Keywords: forgetting, visual working memory, retro-cue, attention, memory decay

It has long been known that covert attention is capable of facilitating visual processing (Helmholtz, 1867). In their seminal studies, Posner and Cohen (1984) showed that presenting a cue at the location of a target facilitates its detection. Subsequent studies have shown that not only is the speed of visual processing influenced by attention, but also the amount of information participants are later capable of reporting. The influence of attention cues on subsequent retrieval from visual working memory (VWM) has typically been studied using *precue* paradigms in which a cue is presented before or during a stimulus array. For instance, Palmer found VWM performance to be sensitive to the number of precued items rather than to the number of items actually displayed on the screen (Palmer, 1990).

Most interestingly, recent findings have shown that attention not only facilitates encoding into VWM (as suggested by the

precue literature), but can also affect information already stored in working memory. Even when applied long after a stimulus array has been removed, such retro-cues can improve performance (Astle, Scerif, Kuo, & Nobre, 2009; Astle, Summerfield, Griffin, & Nobre, 2012; Berryhill, Richmond, Shay, & Olson, 2012; Griffin & Nobre, 2003; Landman, Spekreijse, & Lamme, 2003; Lepsien & Nobre, 2006; Lepsien & Nobre, 2007; Lepsien, Thornton, & Nobre, 2011; Makovski & Jiang, 2007; Makovski, Sussman, & Jiang, 2008; Matsukura, Luck, & Vecera, 2007; Sligte, Scholte, & Lamme, 2008; Sligte, Vandembroucke, Scholte, & Lamme, 2010). All these studies employed a change detection paradigm in which a change between two sequential presentations of stimuli had to be detected. When a cue was directed to the location of a selected item of the first array, up to 10 seconds after the array was extinguished (Astle et al., 2012), it led to a robust advantage in detection performance, relative to trials without cues or those in which a cue was presented with the second array (postcue). Interestingly, subjects cannot completely ignore uninformative retro-cues, which leads to the conclusion that attending to retro-cues is not fully under strategic control (Berryhill et al., 2012).

Why Is Retro-Cueing Advantageous?

These findings raise an interesting question. Unlike in precuing tasks, no additional information can be extracted as a result of a retro-cue, so how does it lead to an advantage in performance? Previous studies have shown that the benefit in performance is not due to grouping processes (Sligte et al., 2008), speed-accuracy trade-offs (Griffin & Nobre, 2003; Lepsien, Griffin, Devlin, & Nobre, 2005), response biases (Griffin & Nobre, 2003), eye movements (Griffin & Nobre, 2003;

Yoni Pertzov, Institute of Neurology and Institute of Cognitive Neuroscience, University College London, Queen Square, London, UK; Paul M. Bays, Sobell Department of Motor Neuroscience and Movement Disorders, Institute of Neurology, University College London; Sabine Joseph, Institute of Neurology and Institute of Cognitive Neuroscience, University College London; Masud Husain, Institute of Neurology and Institute of Cognitive Neuroscience, University College London, and Department of Experimental Psychology and NDCN, University of Oxford, Oxford, UK.

This research was supported by The Wellcome Trust and the NIHR BRC at UCL/UCLH.

Correspondence concerning this article should be addressed to Yoni Pertzov, Institute of Cognitive Neuroscience, University College London, Alexandra House, 17 Queen Square., London, United Kingdom WC1N 3AR. E-mail: y.pertzov@ucl.ac.uk

Matsukura et al., 2007), or articulation (Makovski & Jiang, 2007; Makovski et al., 2008). However, no agreement has been reached regarding the mechanism by which retro-cues enhance performance. Three potential mechanisms, which we discuss in detail below, have been raised in recent debates, but it is important to note that these need not be mutually exclusive.

Prioritization of Item Retrieval

One option is that subjects use the cue to *prioritize* an item when searching through memory. Thus, by initiating their memory search with the cued item during recall, subjects eliminate the need to exhaustively retrieve every item in memory. This view is supported by findings that *valid* retro-cues (directed at the changed object) lead to faster responses but *invalid* cues (directed at one of the nonchanged objects) slow subjects down, because their search presumably initiates from the wrong item. For example, Astle, Summerfield, Griffin, and Nobre (2012) found such differences in response times when retro-cueing arrays of four items, but no difference in the accuracy of report. These results are in line with the prioritization account because uncued items would be equally available for reinspection but just retrieved later: "... a pure prioritisation account, in which cued and uncued items are equally available at probe onset, would predict cue validity effects on RTs but not on accuracy" (Astle et al., 2012). Thus, the prioritization account predicts only effects on response time, but not on the *quality* of report. It is important to note though that Astle et al. (2012) have found evidence for the prioritization account only at low memory loads; high memory loads appear to be associated with the employment of one of the other mechanisms described below, raising the possibility that different strategies might be used in different situations.

Decreased Interference From Subsequent Stimuli

A mechanism that would address differences in accuracy of report is making a retro-cued item *more robust to interference* from, or overwriting by, subsequent new stimuli. Landman et al. (2003) proposed that visual information is initially held in a high capacity but fragile mode (see also Sligte, Lamme, & Scholte, 2006), distinguished from iconic memory by its longer persistence and equivalent sensitivity to isoluminance and high-contrast stimuli (Sligte et al., 2008). This information, it was proposed, is easily interfered upon by subsequent visual stimuli, unless attention has already focused on one of the memorized items. Later studies have added much empirical support for this suggestion, demonstrating that retro-cues can indeed enhance performance by increasing the robustness to interference from subsequent task irrelevant stimuli (Makovski & Jiang, 2007), as well as from the second stimulus array of change detection tasks (Makovski et al., 2008).

Protection From Temporal Degradation

Although the previous account focused on interference originating from stimuli presented *following* the retention interval, another potential mechanism is that retro-cuing leads to protection of the cued item over time *during* the retention interval. In other words, during maintenance, the cued item might be unaffected by *temporal decay* or processes (possibly interference) that degrade other

items in memory. This possibility was originally discarded by Landman et al. (2003) who used a double-retro-cueing paradigm in which a second retro-cue was presented after the first. They reported equivalent performance when the two cues addressed the same compared with different items in memory (see also Lepsien et al., 2011). Their finding suggests that no information was lost from the uncued items when the first cue appeared. Other studies have also failed to observe decreased performance with increasing delay intervals and concluded, therefore, that the protection from degradation account is implausible (Makovski & Jiang, 2007).

However, one study arrived at a very different conclusion. Matsukura et al. (2007) reported a cost when two serially presented retro-cues were directed to different items, supporting the hypothesis that retro-cues operate by avoiding the loss of information within memory. But the drop in performance accuracy was extremely modest (3% of correct detections) and provided only indirect support for this hypothesis. Thus, there is no strong evidence that retro-cues might indeed be able to protect an item from degrading over time. Here we aimed to *directly* test the "protection hypothesis" in more depth using a novel design which potentially provides a more sensitive index of memory report.

All Retro-Cue Studies Have Used a Change Detection Design

As noted earlier, all previous retro-cue studies have used the change detection paradigm in which a change between two successive presentations has to be detected. However, several recent investigations have criticized the use of this task as the gold standard way to probe visual working memory. Detecting a change does not imply perfect recollection of the changed object; nor does detection failure necessarily mean the total absence of memory. For example, detection performance is affected both by the amount of change (Wilken & Ma, 2004) and item complexity (Alvarez & Cavanagh, 2004).

Acknowledging these confounds, one previous retro-cue study introduced an additional test immediately after the change had to be reported (Sligte et al., 2010). These investigators found that only in half the trials on which people were able to detect a change could they also identify the item that changed. Another potential shortcoming of the change blindness design relates to the second presentation of the stimulus array. Because such a presentation can interfere with or overwrite the memory of the first set (Landman, Spekreijse, & Lamme, 2003; Makovski et al., 2008) it makes it difficult to probe memory without any "interference from subsequent stimuli." These considerations suggest that we have to be very cautious when interpreting data from change detection tasks, particularly when using these to index the number of representations maintained in memory. Recent studies in visual working memory have started to use a method that might be more suitable for probing the quality of memory in a sensitive manner.

Using Retro-Cues in a Continuous Report Task

Here we tested the fidelity of memory by using a task that measures quality of recall, or how well an item is remembered, rather than a simple binary measure (e.g., changed or not). This method has been used successfully in examination of capacity limits and resolution of visual working memory (Bays, Catalao, &

Husain, 2009; Bays & Husain, 2008; van den Berg, Shin, Chou, George, & Ma, 2012; Wilken & Ma, 2004). In such tasks participants are required to reproduce from memory an item feature using a *continuous* feature space. Because memory representations are better described as noisy signals rather than all-or-none (present or not) representations (Bays & Husain, 2008; van den Berg et al., 2012; Wilken & Ma, 2004), each response from a continuous feature space is expected to provide much more information (several bits) compared with binary responses (changed vs. unchanged, one bit). Therefore the analog nature of response in this task might make it more appropriate for studying the quality of memory.

Using Variable Delays Between Cue and Probe

Studying the protection account of retro-cuing requires us to gain an understanding of the stability of cued versus uncued working memory representations across time. Although several studies have used variable delays between the stimulus array and the retro-cue (e.g., Slighte et al., 2010), surprisingly no study has systematically manipulated the length of delay between the *retro-cue* and *report* stage. This is all the more remarkable because the protection account generates clear predictions for such a manipulation. Accuracy of the cued item should be stable across time although the accuracy of other items deteriorates. In order to investigate these specific predictions we have used retro-cue designs with variable cue-to-probe intervals as well as analog scale of report. Note that the current study was not designed to distinguish which one of the three hypotheses discussed above is more likely to be correct, but rather to investigate whether retro-cuing might indeed protect an item from degradation over time.

Are Nonspatial Retro-Cues Effective?

In version A of the experiment we used a spatial cue for selecting a single item in memory. However, although *precuing* can direct attention to nonspatial features such as color, it has not been established if retro-cues are also capable of this. One previous study has used object-based retro-cuing (scenes vs. faces; Lepsien & Nobre, 2007) but all other retro-cue studies have reported only on spatial cueing benefits. Therefore, in version B of the experiment, we also investigated whether retro-cues for color lead to performance benefits. If they act to protect an item in the same manner as spatial cues, this would provide important evidence of a general mechanism that operates in working memory.

Method

Experimental Procedures

Twenty-four neurologically normal subjects (age range 19–35 years) participated in version A (12) and B (12) of the experiment after giving informed consent. All subjects reported normal or corrected-to-normal visual acuity. Stimuli were presented at a viewing distance of 60 cm on a 21" CRT monitor.

The experimental designs are illustrated in Figure 1A and Figure 1C for version A and B of the experiment, respectively. Each trial began with the presentation of a central fixation cross (white, 0.8° diameter) for 500 milliseconds. This was followed by a stimulus array consisting of four oriented bars ($2^\circ \times 0.3^\circ$ of visual angle)

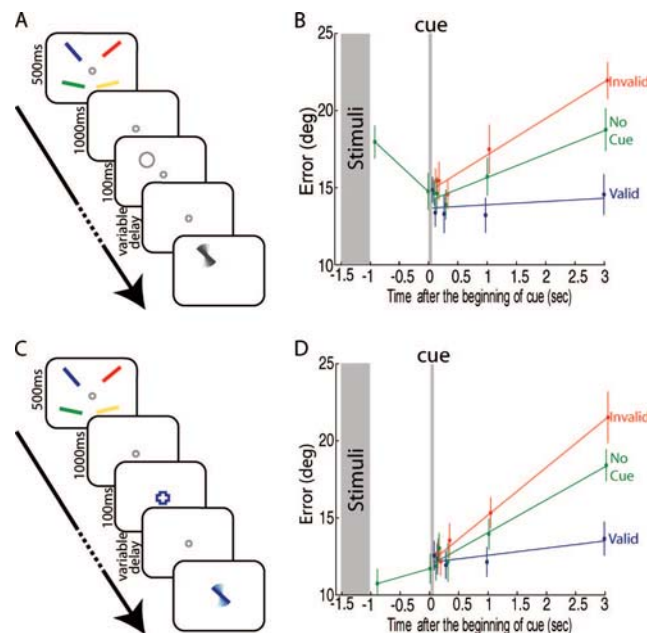


Figure 1. Experimental Design and Results (A) Version A: Spatial cueing. (B) Error in recall for the three different conditions: Blue: Valid, cue matches the probe; Red: Invalid, cue does not match the probe; Green: No Cue, no cue was presented. The linear regression lines for errors following cue onset are plotted in different colors according to the cueing condition. (C) Version B: Color cueing. (D) Results similar to B. Error bars denote SEM across participants ($N = 12$).

presented for 500 ms on a gray background. Each bar had its center arranged at the corners of an $8.8^\circ \times 8.8^\circ$ virtual square centered on the screen. Bars within the same trial differed by at least 10° in orientation, which was otherwise random. The colors of the bars in each trial were randomly selected without repetition out of eight easily distinguishable colors.

Thirty percent of trials had unfilled delays without any retro-cue. These comprised the baseline “no cue” condition. In 70% of trials a retro-cue, indicating the item most likely to be probed, was presented 1 second after the sample stimuli had been extinguished. The cue was either a gray ring at the location of one of the items in the memory array (Version A: location retro-cue), or change in fixation color (Version B: color retro-cue). On 70% of retro-cue trials the cue corresponded to the item that was subsequently probed (*valid condition*). On the remaining 30% of retro-cue trials, one of the other items that had appeared in the sample array was cued (*invalid condition*). In all trials, following a variable delay, a probe item at the same location (Version A) or with the same color (Version B) of one of the items was presented in random orientation.

Subjects adjusted the orientation of the probe using a response dial (Logitech Intl. SA) to match the remembered orientation of the item at the same location (Version A) or with the same color (Version B)—henceforth termed the target. Note that we use the term “target” here simply to distinguish from other items, or nontargets, in the sample array, that were not probed. The probe in no-cue trials was presented at various delays following the stimulus presentation. These delays matched the delays of the cued

trials with the addition of two further time-points 0.1 and 1 second after the stimuli was extinguished.

Eye tracking was performed on half the subjects using a video-based frame-mounted infrared eye tracker (Eye Link1000; SR Research, Canada) with a sampling rate of 1,000 Hz. An online process was used to discard any trials in which participants failed to maintain fixation within 1.6° of the fixation cross (present at all stages except the response). Results were not significantly different between participants who had eye tracking and those who did not.

Each participant performed at least seven blocks of 71 trials (497 trials). Each block included 35 valid and 15 invalid trials, with equal number of trials for each delay duration (seven and three, respectively for 100, 150, 300, 1,000, and 3,000 ms following the cue presentation). The block also included 21 trials without a cue: three trials for each time-point matching the time-points in the cued conditions as well as two additional time-points at 100 and 1,000 ms following the beginning of the delay period.

Analysis

For each trial, a measure of error was obtained by calculating the angular deviation between the orientation reported by the subject and the true orientation of the target item in the sample array. These values were averaged separately for the different trial conditions and durations of delay. To analyze the temporal stability of memory we calculated forgetting slopes using a *linear regression of the error across time* for the five time points (following the cue) for each participant (see Figure 1). Then we performed standard statistical comparisons (paired two-tailed *t* test) between the different cue conditions.

Results

Participants were briefly presented with randomly oriented colored bars and, after a variable delay, were asked to reproduce from memory the orientation of one of the bars, specified either by its location (version A) or color (Version B). In 70% of trials a retro-cue was presented which correctly predicted the probed memory item with 70% validity. Table 1 and Figure 1 display the results from both experiments. Note how even in the baseline no-cue condition (green lines), error (measured with the continuous, analog method) increased rapidly over just 3 seconds.

First we applied a mixed ANOVA with cuing-condition (valid, no-cue, invalid), cue-to-probe delay (100, 150, 300, 1,000, 3,000 ms) as within-subjects factors and version (spatial vs. color cuing) as between subjects factor. This revealed a main effect of cuing-condition, $F(2, 44) = 33$; $p < .0001$, reflecting the significance modulation of performance by the cue. The main effect of delay was also significant, $F(4, 88) = 44$; $p < .0001$, as well as the Cuing-Condition \times Delay interaction, $F(8, 176) = 7$; $p < .0001$, supporting the role of delay in generating a difference in performance across cuing conditions. The version of the experiment (spatial or color retro-cuing) did not generate any significant difference in any effect or interaction ($F_s < 1.7$) suggesting that both types of cuing (spatial and color) lead to similar results.

Next we analyzed how the cue-to-probe delay influenced performance in the different conditions separately for each version. For this purpose we performed linear regression and compared the forgetting slopes. In both versions of the experiment, the fidelity with which the *cued* item was recalled was relatively stable across time (Figure 1: blue lines; mean slope 0.2 deg/sec for spatial and 0.5 deg/sec for color cuing). The gradient of recall accuracy over time was not significantly different from zero [$t(11) < 1.5$; $p > .16$]. In contrast to cued items, the accuracy of neutral, uncued items deteriorated rapidly with time (Figure 1: green lines; mean slope 1.5 deg/sec for spatial and 2.2 for color cuing). These slopes were significantly different from zero [spatial retro-cue $t(11) = 3.2$; $p = .008$; color $t(11) = 9$; $p < .001$] and from the forgetting slopes of the cued items [spatial retro-cue $t(11) = 2.3$; $p = .04$; color $t(11) = 4.8$; $p < .001$].

Forgetting slopes for invalidly cued items (retro-cue directed to items that were not probed at test) were steeper still (Figure 1: red line; mean slope 2.39 deg/sec for spatial and 3.14 for color cuing), with significantly worse performance for invalidly cued items relative to uncued items after 3 seconds [spatial $t(11) = 2.5$; $p = .03$; color $t(11) = 2.3$; $p = .04$]. These findings show that although memory of several items gradually degrades, retro-cues make recall accuracy of selected items relatively stable across time, but importantly at a cost to the accuracy of recall for the other items.

We also quantified the time it takes retro-cues to become active, and thereby lead to a significant difference in performance. In both experimental versions the effect of cue was not observed 300 milliseconds after the cue [Valid vs. Invalid: spatial $t(11) = 1.7$; $p = .12$; color retro-cue $t(11) = 1.1$; $p = .31$] but was clearly

Table 1

Averaged Angle of Error (Degrees) for the two Versions of the Experiment for Different Cuing Conditions and Delays. SEM Across Participants is Reported in Brackets

	Delay from		Delay from				
	end of stimuli display		retro-cue				
	100 ms	1,000 ms	100 ms	150 ms	300 ms	1,000 ms	3,000 ms
Version A							
Invalid cue			15.4 (0.8)	15.4 (1.2)	14.6 (0.9)	17.5 (1.6)	21.9 (1.2)
No-cue	18.0 (1.1)	14.8 (1.2)	14.6 (0.8)	14.6 (1.1)	13.9 (1.0)	15.7 (1.2)	18.7 (1.4)
Valid cue			14.9 (0.8)	13.4 (0.9)	13.3 (1.2)	13.2 (1.1)	14.6 (1.3)
Version B							
Invalid cue			12.6 (1.0)	12.1 (0.9)	13.5 (1.1)	15.3 (1.0)	21.6 (1.7)
No-cue	10.7 (1.0)	11.7 (1.0)	11.8 (0.9)	13.0 (1.0)	12.1 (1.2)	13.9 (1.0)	18.4 (1.1)
Valid cue			12.5 (1.0)	12.3 (1.0)	11.9 (0.9)	12.1 (1.0)	13.6 (1.1)

evident at 1,000 milliseconds [spatial $t(11) = 2.9$; $p = .02$; color retro-cue $t(11) = 2.8$; $p = .02$]. These time constraints are presumably affected by the time it takes to incorporate the cue. However, because of the gradual nature by which a cue might affect performance, even if it is immediately incorporated, some time is expected to elapse before a significant difference in performance might be evident.

Note that participants' performance on no-cue trials at the shortest delays (100 ms following the stimuli) differed between the two experimental versions. In the color retro-cuing paradigm, performance at that time point was better than at all others, as might be expected from the shorter delay. In contrast, larger errors were observed on the spatial-cuing task at 100 ms. This may reflect backward masking: In this task the probe appeared at the same location as the target (but with random orientation) and, therefore, at very short intervals functioned as a visual mask, disrupting processing of the preceding array. Importantly, this effect was abolished by 1 sec, the time at which the retro-cue was presented on cued trials. For all other time points, the dynamics of memory accuracy was similar in both tasks.

Discussion

We studied the fidelity with which visual items are maintained in working memory after a retro-cue selected an item as a probable target. Unlike all previous retro-cue studies, we systematically manipulated the cue-to-probe interval and used an analog scale of report. This method enabled us to directly address a specific hypothesis regarding the mechanism by which retro-cues enhance performance. First, we found that when four items had to be remembered, without any cues, accuracy of recall degraded rapidly over time. However, this decay was almost abolished once a selected item in memory was cued, even though this occurred well after the stimuli were extinguished. But this selective stability came with a cost of enhanced forgetting for other items held in memory. Our results demonstrate that benefits of retro-cuing could come from a dynamic shift of memory resources that protects the temporal stability with which the privileged cued item could be recalled, but at a corresponding cost to the fidelity of recall for uncued items.

The cue-based improvement in temporal stability supports the "protection" hypothesis in which retro-cued items are protected from becoming degraded during the retention interval. Previous investigations using double retro-cuing designs have generated inconclusive results. Three studies have shown that retro-cuing one item and cuing another one later does not in fact lead to worse performance relative to retro-cuing the same item twice (Landman et al., 2003; Lepsien & Nobre, 2007; Lepsien et al., 2011). This suggests that no information is lost from the uncued items after the first cue appears, inconsistent with the protection account. However, two of these studies provided no means of determining whether their observers attended to the item indicated by the first cue on double-cue trials, so their results could be due to observers simply ignoring the first cue and attending exclusively to the second cue.

Lepsien et al. (2007) have used switch/stay cues as the second retro-cue to prevent participants from ignoring the first retro-cue, but still no significant difference in accuracy following stay and switch retro-cues was found. Another study (Matsukura et al.,

2007) which validated that subjects attended to the first cue, did find a cost when applying double retro-cues to two different items, supporting the protection hypothesis. However, the cost was modest (3% of detection rate) and provided only indirect support for this hypothesis. On the other hand, our novel design produced strong and direct support for the "protection hypothesis." Three seconds following an invalid cue, errors were almost twice the magnitude compared with following a valid cue. This strong retro-cuing effect occurred because performance of validly cued items did not degrade across the time, providing the first direct support for the protection hypothesis.

Unifying Account: Protection From Interference

It is important to note that our findings regarding the changes in temporal stability of memory recollection do not argue against other mechanisms that might also be responsible for the retro-cuing effect. The possibilities discussed previously—prioritization, rendering items robust to interference and protection from degradation—are certainly not mutually exclusive. On the contrary, we believe that our findings could be reconciled with other explanations and perhaps generate a more unifying mechanism. As described in the beginning of this article, a growing number of studies have demonstrated that a retro-cue makes items more resilient to interference caused by subsequently presented stimuli such as the test array (Landman et al., 2003; Makovski & Jiang, 2007; Makovski et al., 2008; Sligte et al., 2008, 2010).

The current findings may extend this view by suggesting that retro-cued items are protected from degradation caused by interference between items residing simultaneously in memory. Single items can be maintained with only minor loss across time but multiple memory items may compete for memory resources and suppress each other's representation leading to memory degradation (Bahcall & Kowler, 1999; Edin et al., 2009; Oberauer, 2009). Directing attention to a single memory item can counteract the interference from other memory items, enhancing the relative benefit in which selected items are recalled. Therefore, a parsimonious and unifying explanation for both findings would be that retro-cuing protects selected items in memory from interference applied either by novel incoming stimuli or by other items residing in memory, over time.

Invalid Cues

But this protection does not come without a cost. Our results clearly show that the retro-cue makes the nonselected items even more fragile than items in trials without any cue. Note that decreased accuracy for the invalid and no-cue conditions might arise from several sources. One possibility is decreased precision of target representations. Alternatively, such items might even be completely lost from memory, which would be expected to lead to increased probability of guessing (Zhang & Luck, 2009). Increased error might also result from participants systematically misreporting the wrong item in memory (Bays et al., 2009; Gorgoraptis, Catalao, Bays, & Husain, 2011). Distinguishing between these various sources is theoretically possible but crucially requires an analysis of the *distribution* of responses and, therefore, a very large number of responses for every condition. However, by definition "invalid" trials are

scarce (a high proportion of “invalid” trials would make them “valid”), making any model fit very noisy. Distinguishing the sources of errors might be possible using designs that are adequately powered to examine this issue, but the key aim of the current study was to determine whether there is any evidence that retro-cues protect an item from degradation over time.

Neural Basis of Degradation in Immediate Recall

Although most forgetting curves at longer delays (minutes to years) describe a curvilinear function (e.g., exponential or power; Della Sala, 2010), under our experimental settings, using short delays and an analog measure of error, error appeared to increase linearly over time (see Figure 1). Memory over several seconds is widely believed to be related to persistent patterns of elevated or inhibited neural activity (e.g., Miller, Erickson, & Desimone, 1996). Because the persistence time of individual neurons is much shorter than the intervals of observed persistent activity, neural models typically suggest that maintained patterns of neural activity, or attractors, are generated and stabilized through reverberating positive feedback (e.g., Wei, Wang, & Wang, 2012). In line-attractor networks, recurrent excitation leads to the formation of a continuous line of stable states enabling the storage of analog values. The existence of stochastic neural noise is expected to lead to a random drift along the stable dimension of the network, and thus, to degradation of the stored memory over time. The time-course of this degradation has been suggested to lead to a linear relationship between the retention interval and the *variance* between the recalled variable and its true value (Burak & Fiete, in press). Hence the linear forgetting slopes we have found with respect to the mean absolute error (rather than variance of error) are somewhat steeper than the slopes expected solely from information degradation as a result of random drift in an isolated line-attractor network. Further studies are needed to assess the hypothesized link between neural-drift and immediate forgetting in continuous report tasks.

Memory Representations Are Not All Equal: Attention to an Item in Working Memory

Our findings have one clear, but nontrivial, consequence: not all working memory representations are equal. One memory representation can be temporally stable—attended to or given a privileged status—although others are more fragile. The special state of the cued representation could theoretically be accomplished by “covert visuospatial rehearsal” (Baddeley, 2007), “attentional refreshing” (Barrouillet, Bernardin, & Camos, 2004; Cowan, 1988), or “visual imagery” (Baddeley, 2007). Thus, our results are consistent with the privileged state of the item under the “focus of attention” in working memory (Jonides et al., 2008; Oberauer, 2002; Olivers, Peters, Houtkamp, & Roelfsema, 2011). The benefit to the retro-cued item—with a corresponding cost to uncued items—echoes findings from precuing studies of visual attention (Posner & Cohen, 1984), suggesting that redeployment of resources to one item in memory leads to loss of resource to others.

Of course, the mechanism that keeps an item in a privileged state might not be identical to attention resources deployed in visual processing. We are not proposing that it has to be. Critically,

the retro-cuing experimental design controls well for visual processing, which happens long before the cue is presented. Thus, the resource redeployment triggered by the retro-cue is consistent with the new taxonomical definition of “internal” or “reflective” attention (Chun, Golomb, & Turk-Browne, 2011; Chun & Johnson, 2011). And indeed evidence exists that the attentional processes triggered by precues might have different properties than attention triggered by retro-cues. Although precuing multiple items leads to better change detection performance for all the selected items, the retro-cue benefit was found to be abolished when more than one item is cued (Makovski & Jiang, 2007; but see Lepsien et al., 2011). Thus, there might be an important distinction between “internal” and “perceptual” attention when regarding retro-cues. The former might be less effective than the latter when directed to multiple items. Further studies would be needed to characterize the specific similarities and differences between the two types of attention process.

The current study nevertheless contributes to the view that resources can be focused on a privileged item in working memory but at a cost to other items: by demonstrating that uncued items in memory undergo deterioration, and, therefore, exist in a temporally fragile mode. Various other differences between retro-cued and uncued memory representation have led to the suggestion that visual short-term memory (STM) should be divided into two types. A high capacity but fragile memory has been proposed to exist in addition to the limited capacity but relatively stable memory store that has traditionally been invoked (Landman et al., 2003; Sligte et al., 2008, 2010). Note, however, that the fragility of memory was previously discussed only in relation to interference from incoming stimuli. Our results show that representations in STM also differ with respect to their temporal stability, even without new items to encode into memory.

Nonspatial Cues

One additional distinction between “internal” and “perceptual” attention might have been the applicability of nonspatial cues. Although nonspatial (e.g., color) precueing can clearly improve performance, evidence for benefits of nonspatial retro-cues is scarce. Berryhill, Richmond, Shay, and Olson (2012) have used four digit cues that map to four different items in the memory array. Despite such cues being fully 100% predictive of the future changed item, no performance benefits were observed, leading to the conclusion that retro-cue benefits are not sensitive to such cues. However, the failure to incorporate this cue could be related to several experimental details, such as the complex mapping (digits to space) between cues and items as well as to the relatively short interval allotted for this mapping (400 ms).

Only one study has reported a nonspatial, retro-cuing effect. Lepsien and Nobre (2007) have shown that cuing the category (faces or scenes) of previously presented stimuli can lead to specific benefits in remembering the stimuli belonging to this category. The aim of our experimental version B was to investigate how robust our results were to various manipulations and specifically to check if nonspatial cues can be incorporated to generate performance benefits. The current study is the first to show that subjects are able to incorporate color retro-cues in order to direct attention within working memory. This is not trivial, especially because working memory seems to preserve the spatial properties

of the stimuli it represents (e.g., Griffin & Nobre, 2003). Therefore, spatial location might have played a special role in accessing working memory. In fact, subjects might still have used a spatial strategy even in the color cuing version. For instance, when a specific color-cue appears, subjects might have allocated their spatial attention to the location of the item with that color. Thus, they could essentially recode the cue into spatial terms and apply a spatial attention mechanism. The putative indirect effect of color cuing does not seem to be reflected in the time it takes to incorporate the cue because the two experimental versions lead to comparable delays before observing a significant retro-cue effect. Nevertheless, this account is feasible and deserves examination in future studies. In any event, the effects of color retro-cues virtually replicated our findings from spatial retro-cues, demonstrating that the cuing effects on temporal stability are robust and generalize across experimental differences.

In conclusion, both spatial and color retro-cues are capable of protecting a selected item in memory from gradual degradation and keeping it in a privileged, “attended” status. However, such protection comes with a cost of enhanced, rapid forgetting of other items in memory.

References

- Alvarez, G. A., & Cavanagh, P. (2004). The capacity of visual short-term memory is set both by visual information load and by number of objects. *Psychological Science, 15*, 106–111. doi:10.1111/j.0963-7214.2004.01502006.x
- Astle, D. E., Scerif, G., Kuo, B. C., & Nobre, A. C. (2009). Spatial selection of features within perceived and remembered objects. *Frontiers in Human Neuroscience, 3*, 6.
- Astle, D. E., Summerfield, J., Griffin, I., & Nobre, A. C. (2012). Orienting attention to locations in mental representations. *Attention, Perception, & Psychophysics, 74*, 146–162.
- Baddeley, A. D. (2007). *Working memory, thought, and action*. Oxford, UK: Oxford University Press. doi:10.1093/acprof:oso/9780198528012.001.0001
- Bahcall, D. O., & Kowler, E. (1999). Attentional interference at small spatial separations. *Vision Research, 39*, 71–86. doi:10.1016/S0042-6989(98)00090-X
- Barrouillet, P., Bernardin, S., & Camos, V. (2004). Time constraints and resource sharing in adults’ working memory spans. *Journal of Experimental Psychology: General, 133*, 83–100. doi:10.1037/0096-3445.133.1.83
- Bays, P. M., Catalao, R. F., & Husain, M. (2009). The precision of visual working memory is set by allocation of a shared resource. *Journal of Vision, 9*, 1–11. doi:10.1167/9.10.7
- Bays, P. M., & Husain, M. (2008). Dynamic shifts of limited working memory resources in human vision. *Science, 321*, 851–854. doi:10.1126/science.1158023
- Berryhill, M. E., Richmond, L. L., Shay, C. S., & Olson, I. R. (2012). Shifting attention among working memory representations: Testing cue type, awareness, and strategic control. *The Quarterly Journal of Experimental Psychology, 65*, 426–438. doi:10.1080/17470218.2011.604786
- Burak, Y., & Fiete, I. R. (in press). Fundamental limits on persistent activity in networks of noisy neurons. *Proceedings of the National Academy of Sciences*.
- Chun, M. M., Golomb, J. D., & Turk-Browne, N. B. (2011). A Taxonomy of External and Internal Attention. *Annual Review of Psychology, 62*, 73–101.
- Chun, M. M., & Johnson, M. K. (2011). Memory: Enduring traces of perceptual and reflective attention. *Neuron, 72*, 520–535. doi:10.1016/j.neuron.2011.10.026
- Cowan, N. (1988). Evolving conceptions of memory storage, selective attention, and their mutual constraints within the human information-processing system. *Psychological Bulletin, 104*, 163–191. doi:10.1037/0033-2909.104.2.163
- Della Sala, S. (2010). *Forgetting*. New York, NY: Psychology Press.
- Edin, F., Klingberg, T., Johansson, P., McNab, F., Tegnér, J., & Compte, A. (2009). Mechanism for top-down control of working memory capacity. *Proceedings of the National Academy of Sciences, 106*, 6802–6807. doi:10.1073/pnas.0901894106
- Gorgoraptis, N., Catalao, R. F., Bays, P. M., & Husain, M. (2011). Dynamic updating of working memory resources for visual objects. *The Journal of Neuroscience, 31*, 8502–8511. doi:10.1523/JNEUROSCI.0208-11.2011
- Griffin, I. C., & Nobre, A. C. (2003). Orienting attention to locations in internal representations. *Journal of Cognitive Neuroscience, 15*, 1176–1194. doi:10.1162/089892903322598139
- Helmholtz, H. (1867). *Treatise on physiological optics, Vol. III*. (Trans. and Ed., JPC Southall.) New York, NY: Dover. (Translated from the 3rd German edition, English edition 1962)
- Jonides, J., Lewis, R. L., Nee, D. E., Lustig, C. A., Berman, M. G., & Moore, K. S. (2008). The mind and brain of short-term memory. *Annual Review of Psychology, 59*, 193–224. doi:10.1146/annurev.psych.59.103006.093615
- Landman, R., Spekreijse, H., & Lamme, V. A. (2003). Large capacity storage of integrated objects before change blindness. *Vision Research, 43*, 149–164. doi:10.1016/S0042-6989(02)00402-9
- Lepsien, J., Griffin, I. C., Devlin, J. T., & Nobre, A. C. (2005). Directing spatial attention in mental representations: Interactions between attentional orienting and working-memory load. *Neuroimage, 26*, 733–743. doi:10.1016/j.neuroimage.2005.02.026
- Lepsien, J., & Nobre, A. C. (2006). Cognitive control of attention in the human brain: Insights from orienting attention to mental representations. *Brain Research, 1105*, 20–31. doi:10.1016/j.brainres.2006.03.033
- Lepsien, J., & Nobre, A. C. (2007). Attentional modulation of object representations in working memory. *Cerebral Cortex, 17*, 2072. doi:10.1093/cercor/bhl116
- Lepsien, J., Thornton, I., & Nobre, A. C. (2011). Modulation of working-memory maintenance by directed attention. *Neuropsychologia, 49*, 1569–1577. doi:10.1016/j.neuropsychologia.2011.03.011
- Makovski, T., & Jiang, Y. V. (2007). Distributing versus focusing attention in visual short-term memory. *Psychonomic Bulletin & Review, 14*, 1072–1078. doi:10.3758/BF03193093
- Makovski, T., Sussman, R., & Jiang, Y. V. (2008). Orienting attention in visual working memory reduces interference from memory probes. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 34*, 369. doi:10.1037/0278-7393.34.2.369
- Matsukura, M., Luck, S. J., & Vecera, S. P. (2007). Attention effects during visual short-term memory maintenance: Protection or prioritization? *Attention, Perception, & Psychophysics, 69*, 1422–1434.
- Miller, E. K., Erickson, C. A., & Desimone, R. (1996). Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *The Journal of Neuroscience, 16*, 5154–5167.
- Oberauer, K. (2002). Access to information in working memory: Exploring the focus of attention. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 28*, 411–421. doi:10.1037/0278-7393.28.3.411
- Oberauer, K. (2009). Design for a working memory. *Psychology of learning and motivation, 51*, 45–100. doi:10.1016/S0079-7421(09)51002-X
- Olivers, C. N., Peters, J., Houtkamp, R., & Roelfsema, P. R. (2011). Different states in visual working memory: When it guides attention and when it does not. *Trends in Cognitive Sciences, 15*, 327–334. doi:10.1016/j.tics.2011.05.004
- Palmer, J. (1990). Attentional limits on the perception and memory of visual information. *Journal of Experimental Psychology: Human Per-*

- ception and Performance*, 16, 332–350. doi:10.1037/0096-1523.16.2.332
- Posner, M. I., & Cohen, Y. (1984). *Components of visual orienting*. Hillsdale, NJ: Erlbaum.
- Sligte, I. G., Lamme, V. A., & Scholte, H. S. (2006). Iconic memory revisited: A plea for a distinction between a retinal and cortical icon. *Journal of Vision*, 6, 989. doi:10.1167/6.6.989
- Sligte, I. G., Scholte, H. S., & Lamme, V. A. (2008). Are there multiple visual short-term memory stores? *PLoS One*, 3, e1699.
- Sligte, I. G., Vandenbroucke, A. R., Scholte, H. S., & Lamme, V. A. (2010). Detailed sensory memory, sloppy working memory. *Frontiers in Psychology*, 1, 175.
- van den Berg, R., Shin, H., Chou, W. C., George, R., & Ma, W. J. (2012). Variability in encoding precision accounts for visual short-term memory limitations. *Proceedings of the National Academy of Science*, 109, 8780–8785.
- Wei, Z., Wang, X. J., & Wang, D. H. (2012). From distributed resources to limited slots in multiple-item working memory: A spiking network model with normalization. *The Journal of Neuroscience*, 32, 11228–11240. doi:10.1523/JNEUROSCI.0735-12.2012
- Wilken, P., & Ma, W. (2004). A detection theory account of change detection. *Journal of Vision*, 4, 1120–1135. doi:10.1167/4.12.11
- Zhang, W., & Luck, S. J. (2009). Sudden death and gradual decay in visual working memory. *Psychological Science*, 20, 423. doi:10.1111/j.1467-9280.2009.02322.x

Received June 25, 2012

Revision received October 5, 2012

Accepted October 15, 2012 ■