The benefit of forgetting

Melonie Williams, Sang W. Hong, Min-Suk Kang, Nancy B. Carlisle & Geoffrey F. Woodman

Psychonomic Bulletin & Review

ISSN 1069-9384

Psychon Bull Rev
DOI 10.3758/s13423-012-0354-3

Springer

Psychon Bull Rev
DOI 10.3758/s13423-012-0354-3

BRIEF REPORT

The benefit of forgetting

Melonie Williams · Sang W. Hong · Min-Suk Kang · Nancy B. Carlisle · Geoffrey F. Woodman

Abstract Recent research using change-detection tasks has shown that a directed-forgetting cue, indicating that a subset of the information stored in memory can be forgotten, significantly benefits the other information stored in visual working memory. How do these directed-forgetting cues aid the memory representations that are retained? We addressed this question in the present study by using a recall paradigm to measure the nature of the retained memory representations. Our results demonstrated that a directed-forgetting cue leads to higher-fidelity representations of the remaining items and a lower probability of dropping these representations from memory. Next, we showed that this is made possible by the to-be-forgotten item being expelled from visual working memory following the cue, allowing maintenance mechanisms to be focused on only the items that remain in visual working memory. Thus, the present findings show that cues to forget benefit the remaining information in visual working memory by fundamentally improving their quality relative to conditions in which just as many items are encoded but no cue is provided.

M. Williams · M.-S. Kang · G. F. Woodman
Vanderbilt University, Nashville, TN, USA

S. W. Hong
Florida Atlantic University, Boca Raton, FL, USA

M.-S. Kang
Sungkyunkwan University, Seoul, South Korea

N. B. Carlisle
University of California, Davis, CA, USA

G. F. Woodman (✉)
Department of Psychology, Vanderbilt University, PMB 407817,
2301 Vanderbilt Place,
Nashville, TN 37240-7817, USA
e-mail: geoffrey.f.woodman@vanderbilt.edu

Keywords Visual working memory · Human memory · Cognitive and attentional control

Research on directed forgetting in long-term memory has supported the claim that we can volitionally select which representations are retained in memory (Anderson & Green, 2001; Anderson & Levy, 2009; Johnson, 1994; MacLeod, 1999). When participants are instructed to forget items on a list, recall is significantly lower for those words than for other words on the studied list. Williams and Woodman (2012) recently examined how directed-forgetting cues change what is maintained in visual working memory. In that study, participants were always shown six colored squares to remember. On a subset of the trials, they were provided a cue during the retention interval indicating that they could forget three of the six squares (i.e., a directed-forgetting cue). These cues resulted in superior performance at detecting changes, relative to baseline trials without a cue. Furthermore, evidence from an event-related potential (ERP) experiment indicated that directed-forgetting cues improve change detection for the remaining information in working memory by denying task-irrelevant information access to limited-capacity maintenance processes. This denial of access allows maintenance mechanisms to be focused on the remaining representations in visual working memory. If this explanation is correct, then focusing maintenance on the task-relevant representations in visual working memory could either (1) increase their fidelity relative to conditions in which a set of memoranda are shown and no cues about future relevance are shown, or (2) prevent the remaining representations from being dropped from visual working memory (Zhang & Luck, 2009). In the present study, we sought to determine the source of the benefits conferred by cues to forget information in visual working memory.

To determine how directed-forgetting cues benefit the representations in visual working memory that remain task

Published online: 01 December 2012

relevant, we needed to measure more detailed aspects of the nature of the representations than the accuracy-based measures provided by change-detection tasks. Recently, researchers have developed techniques for measuring the precision and number of representations held in visual working memory (Bays & Husain, 2008; Zhang & Luck, 2008). For example, Zhang and Luck (2008) showed how mixture modeling could be used in conjunction with a recall paradigm to simultaneously measure the number and resolution of representations that are stored. Using this approach, they turned to the question of what happens to items stored in visual working memory across time; do the items gradually decay, or are the representations suddenly lost (i.e., sudden death; Zhang & Luck, 2009)? Their findings showed that information is lost from visual working memory because items are dropped or suddenly die rather than gradually decaying. In the present study, we used the same methodological approach (i.e., a recall task and mixture modeling) to determine whether a cue instructing participants to forget some information would benefit the remaining representations. This could come about either by increasing the resolution and fidelity of the remaining items, because maintenance mechanisms could then be focused on them, relative to conditions in which this cognitive resource was spread across more items, or by preventing their sudden death and total loss. Finally, it is possible that focusing maintenance relative to spreading it more thinly across more items might have both of these effects.

Experiment 1

The goal of Experiment 1 was to measure the nature of memory representations following a cue to forget other items. We used a cued-recall procedure similar to that of Zhang and Luck (2008). Observers were asked to remember one or two colored squares. As is illustrated in Fig. 1, the stimuli were presented briefly, followed by a short retention interval, and then a color wheel with all possible colors and with a box around one location marking the item that needed to be reported. Three trial types occurred during this experiment. On *one-object trials*, we presented only one colored square in the memory array. On *two-object trials*, two colored squares were presented during the memory array. Finally, on *two-cued trials*, two colored squares were presented in the memory array, and a cue to forget one square appeared during the retention interval. During test, only one square was probed on all three types of trials, and observers made a response by clicking on the color wheel.

If focusing maintenance following a cue to forget some information benefits the remaining representations by improving their fidelity, then we should observe better

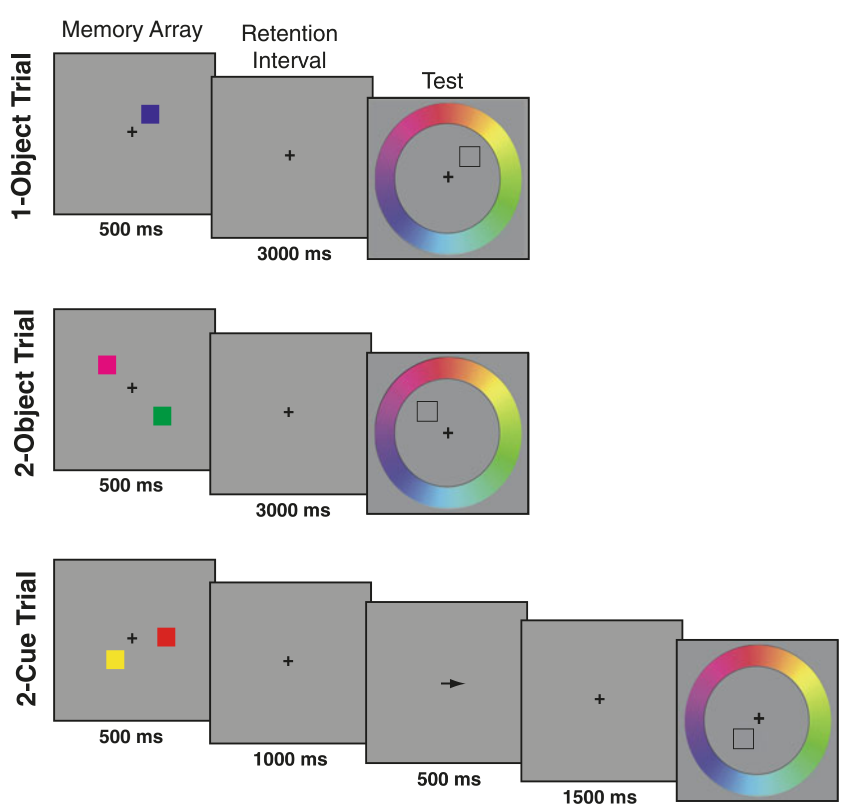

Fig. 1 Example stimuli and the three trial types in Experiment 1. The top panel shows a one-object trial, the middle panel a two-object trial, and the bottom panel a two-cued trial

precision for the tested representation during two-cued as compared to two-object trials. This would be observed as a smaller standard deviation (σ) of the distribution of responses around the actual color of the tested object. However, if cues to forget prevent the sudden death of the remaining memory representations that receive the full benefit of focused maintenance, then the probability of having retained the tested item should be higher on the two-cued trials relative to the two-object baseline. This would be observed as a higher *Pm* (probability of memory) due to the presence of fewer pure guesses that were distant in color space from the color actually shown in the memory array. Finally, it is possible that a cue to forget benefits the remaining representations in both of these ways, improving their fidelity and probability of being retained in visual working memory across the entire retention interval.

Method

Participants A group of 20 volunteers 18–35 years old reported normal or corrected-normal vision and provided informed consent prior to the experiment.

Stimuli and procedure We presented stimuli using MATLAB (The MathWorks, Natick, MA) and the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997) on a calibrated CRT display (Sony Trinitron Multiscan E540) in a near dark room, so that the display was the predominant light source. The spectral power distributions of the red, green, and blue phosphors were measured using a spectroradiometer (Ocean Optics, Model USB4000). The relative light level of each gun at every digital value ($256 = 2^8$ levels) was measured with a Minolta colorimeter (Model No. CA-100). The effective viewing distance was 90 cm.

The stimuli were either one or two colored squares (each $0.64 \times 0.64°$ of visual angle) presented on a gray background (37.7 cd/m^2). The colors of the stimuli were randomly chosen from 180 colors, defined within CIE-Lab color space. The stimulus locations were randomly selected from 12 equally spaced positions on a virtual circle whose radius was 2.5° from the central fixation. For two-object trials, both locations were selected randomly with at least a 90° angle between them. The test stimulus was a color wheel based on that of Zhang and Luck (2008). The color wheel was composed of 180 segments with invisible boundaries, each of which contained a color from CIE-Lab space. The colors of the wheel were centered at ($L = 60$, $a = 0$, $b = 20$) with a radius of 65. The thickness of the wheel was 1.78°, with a 5.22° inner radius. The color wheel randomly rotated across trials to prevent specific color locations being predictable during the experiment. The participants performed a concurrent articulatory suppression task in order to minimize verbal recoding of the object colors. The stimuli were the letter strings "a, b, c, d," "w, x, y, z," "1, 2, 3, 4," and "6, 7, 8, 9," centered at fixation. The forget cue stimuli were black arrows (approximately $0.64 \times 0.32°$) presented at the center of the display, briefly replacing the fixation point.

Each one-object and two-object trial began with a memory array presented for 500 ms, followed by a 3,000-ms blank retention interval. On two-cued trials, the forget cue appeared for a duration of 500 ms after 1,000 ms had passed from the offset of the memory items, followed by the rest of the retention interval, which lasted 1,500 ms. At test, a color wheel appeared along with an outlined square centered in the location of a previously presented square, indicating that participants should report the color of that square. The participants were instructed to use the mouse to click the best-matching color on the wheel. To stress accuracy, the responses were not timed. Feedback regarding color choices was provided immediately after each response, by using an outlined circle on the color wheel to mark the actual color value.

All trial types (one-object, two-object, and two-cued) were randomly intermixed, with each type occurring 100 times across six blocks. At the start of a block, participants were given the articulatory suppression string and told to repeat the letters nonstop. The letter sequence changed at the start of each block. After every 50 trials, the observers were allowed to take a short break before pressing the spacebar to continue.

Data analysis Due to the circular nature of the stimulus space, we analyzed participants' responses using a quantitative model that fitted a von Mises distribution. Using these parameters, we were able to calculate Kappa (κ) and *Pm*. Variability was measured as the width under the Gaussian component of the model and reflected the precision of the memory trace. Kappa—calculated using the equation $f(x|\mu,\kappa) = e^{\kappa\cos(y-\mu)}/[2\pi I_0(\kappa)]$ —was then translated to σ, with $1/\kappa$ being equal to σ^2.

One minus the probability of failure, measured as the height of the uniform distribution component of the mixture model, was then calculated to derive *Pm*, which reflects the probability that the representation was in memory at the time of test.

Results

Figure 2a shows histograms of the responses across the three trial types as well as the mixture model fits. The *Pm* and σ values of the mixture modeling analyses are shown in Fig. 2b and c. The effect of cuing participants to forget one of the two items on the two-cued trials increased *Pm* and decreased σ to levels approaching those when only one item was encoded and maintained on the one-object trials. Differences in *Pm* and σ were analyzed using analyses of variance (ANOVAs) with the within-subjects factor Trial Type (one-object, two-object,

Fig. 2 Results from Experiment 1. **a** Histogram of all responses (in degrees from the actual color value in color space), separated by trial type. Each line represents the mixture distribution's best fit to the data. The inset box shows the effect of trial type on the tails. **b** *Pm* values (± 2 *SEMs*) and **c** *σ* values (± 2 *SEMs*) derived from the mixture model analyses

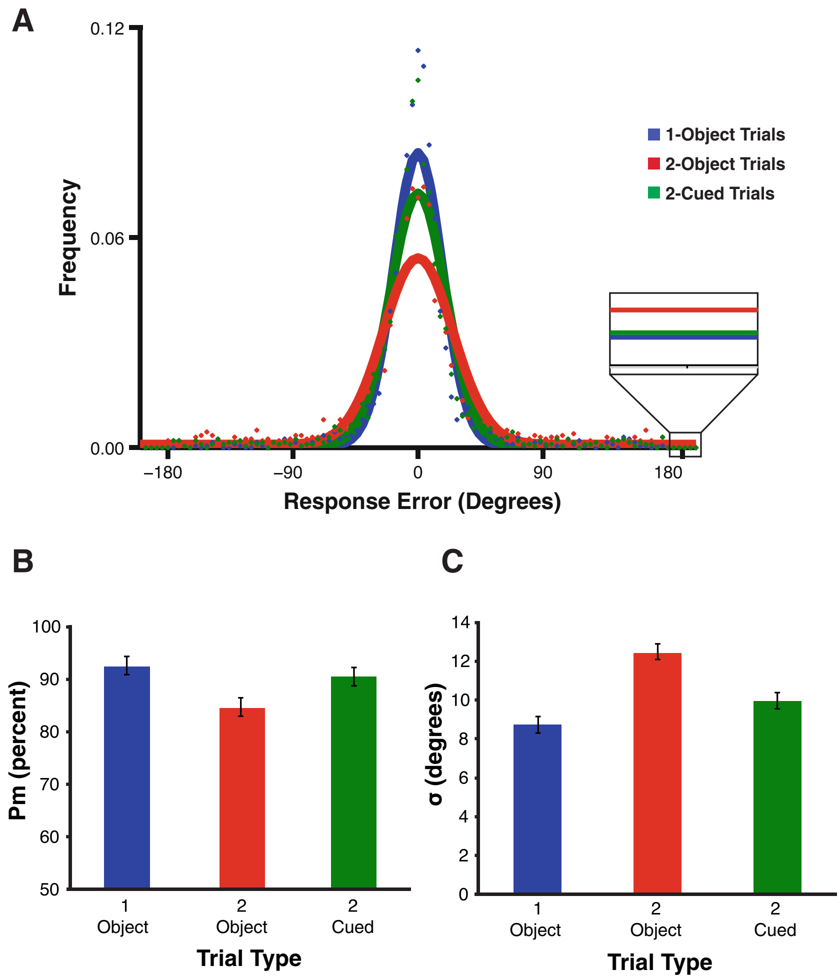

and two-cued). We found significant main effects of trial type on both *Pm*, $F(2, 38) = 13.80$, $p < .001$, and σ, $F(2, 38) = 29.77$, $p < .001$.

Planned comparisons of the two-object trials and two-cued trials showed significant differences in both *Pm*, $t(19) = 3.46$, $p < .05$, and σ, $t(19) = 6.63$, $p < .001$. When comparing one-object to two-object trials, both *Pm*, $t(19) = 5.19$, $p < .001$, and σ, $t(19) = 6.88$, $p < .001$, were significantly different. The comparisons of one-object with two-cued trials only found a significant difference in σ, $t(19) = 2.77$, $p < .05$; the values of *Pm* for these trial types were not significantly different, $t(19) = 1.49$, $p = .15$, showing that the cues to forget one item can make memory for the remaining item similar to the response metrics in the situation in which only one item was ever stored.

Discussion

The findings of Experiment 1 support the hypothesis that focusing maintenance by cuing an observer to forget an item benefits the remaining representation in visual working memory in two ways: It can increase the fidelity of the remaining representation and prevent its loss. The σ findings provided evidence of a boost in precision for the item that people were not cued to forget; the *Pm* findings showed that these cues also prevented the other item from being lost from visual working memory across the remaining 1.5-s retention interval. The difference in standard deviations between the two-cued trials and the one-object trials suggests that some amount of information loss did occur during the second prior to the cue presentation, before maintenance could be focused on just the task-relevant item. What is still unclear is what happened to the item that people were cued to forget. Was that item discarded from working memory, or was it left to simply to passively decay or die from neglect in the visual working memory store? Experiment 2 was designed to distinguish between these possibilities.

Experiment 2

We had two goals in Experiment 2. The first was to see whether the findings from Experiment 1 would replicate with a new and larger group of observers; the second was to observe what happens to the representation of the item that was made task-irrelevant by a cue to forget. It is possible that the benefits observed in Experiment 1 were due to the participants actively expelling the cued item from visual working memory, or they might have simply failed to actively maintain the item that the cue had indicated should be forgotten. To distinguish between these hypotheses, we needed to probe memory for the item that we had cued people to forget. This could be difficult, because once we began testing memory for items that the cues indicated would not be tested, participants might begin ignoring the cues (see Williams & Woodman, 2012). Therefore, we wanted to probe memory infrequently for these forgettable items and to focus on the first such memory probes, when people would not have become sensitized to the possible task-relevance of these memories. This meant that we needed to collect data from a large number of participants to have enough observations to feed the mixture modeling analyses.

The design of Experiment 2 was a shorter version of Experiment 1, with the addition of invalid trials. Unbeknownst to the participants, on five of the two-cued trials the cue was invalid, because the to-be-forgotten item would be probed at test. If the forget cues were responded to by simply not actively maintaining the cued item, then we should see evidence that it was allowed simply to passively decay over time. This would mean that when the participants' memories were probed for the to-be-forgotten object, we would observe precision and *Pm* values similar in form to those in the two-object trials, simply further reduced due to the absence of active maintenance of this visual working memory representation. Note that a probabilistic sudden-death hypothesis (Zhang & Luck, 2009), in which the likelihood of death increases as time passes, would make similar predictions. However, if after a cue to forget an item, the item were expelled from working memory, then we should see that responses to these surprise memory probes resulted in random guessing about the color of this item. That is, we should observe low *Pm* for those items, with no Gaussian peak in the distribution around the color initially presented (i.e., a flat line).

Method

Participants We collected data from 100 new volunteers from the same pool after informed consent was obtained.

Stimuli and procedure The stimuli and procedures were similar to those of Experiment 1, with the addition of the infrequent invalid trials. Because we were primarily interested in participants' responses following the first invalid cue, we shortened the experiment to 50 trials in total: 15 one-object trials, 15 two-object trials, 15 two-cued trials, and five invalid trials following directed-forgetting cues. The participants were given the same instructions as in Experiment 1 and were not warned about the misleading, invalid cues. On trials in which an invalid cue appeared, the participants were directed to ignore that item but then probed on the item during recall.

Results

Figure 3a and b shows that we replicated the pattern of results from Experiment 1 on the trial types that were the same in Experiment 2, with slightly higher overall performance in Experiment 2 than in Experiment 1. However, a cue to forget one of the two items in memory resulted in a higher *Pm* and a lower σ, as compared to the two-object trials without a directed-forgetting cue during the retention interval. In contrast to the high *Pm* and low σ for one-object, two-object, and two-cued trials in Experiments 1 and 2, the invalid trials of Experiment 2 yielded little evidence that the item that the cue had indicated could be forgotten was in memory, with striking variability of the responses and little or no clustering of responses around the actual color of the item that participants had been cued to forget.

The within-subjects ANOVA supported the observation that the original conditions used in Experiment 1 exhibited the same pattern that we saw in Experiment 2. We found significant main effects of trial type (one-object, two-object, and two-cued trials) on *Pm*, $F(2, 198) = 4.42$, $p < .05$, and σ, $F(2, 198) = 26.20$, $p < .001$. As in Experiment 1, our planned comparisons showed that across the two-object and two-cued trials, the difference between the *Pm*s approached significance, $t(99) = 1.85$, $p = .07$, and we found a significant difference in σ, $t(99) = 2.8$, $p < .001$. When comparing the one-object and two-object trials, both *Pm*, $t(99) = 2.72$, $p < .01$, and σ, $t(99) = 6.95$, $p < .001$, were significantly different. The comparisons of the one-object and two-cued trials only revealed that σ, $t(99) = 4.56$, $p < .01$, differed significantly across these trial types, as the *Pm* values for these trial types were not statistically distinguishable, $t(99) = 1.06$, $p = .29$. These analyses mirrored the essential findings reported in Experiment 1. Statistical support for this conclusion came from entering *Pm* and σ data into 2 × 3 ANOVAs with the between-subjects factor Experiment (1 vs. 2) and the within-subjects factor Trial Type (one-object, two-object, and two-cued). Neither ANOVA yielded a significant Experiment × Trial Type interaction, *F*s < 1.0, due to the similar patterns across the dependent variables in both experiments. Both yielded significant effects of trial type [*Pm*, $F(2, 236) = 9.62$, $p < .001$; σ, $F(2, 236) = 19.12$, $p < .001$], and the effect of experiment was significant in the *Pm* ANOVA, $F(1, 118) = 27.00$, $p < .001$, due to higher probabilities of remembering across all trial types in Experiment 2.

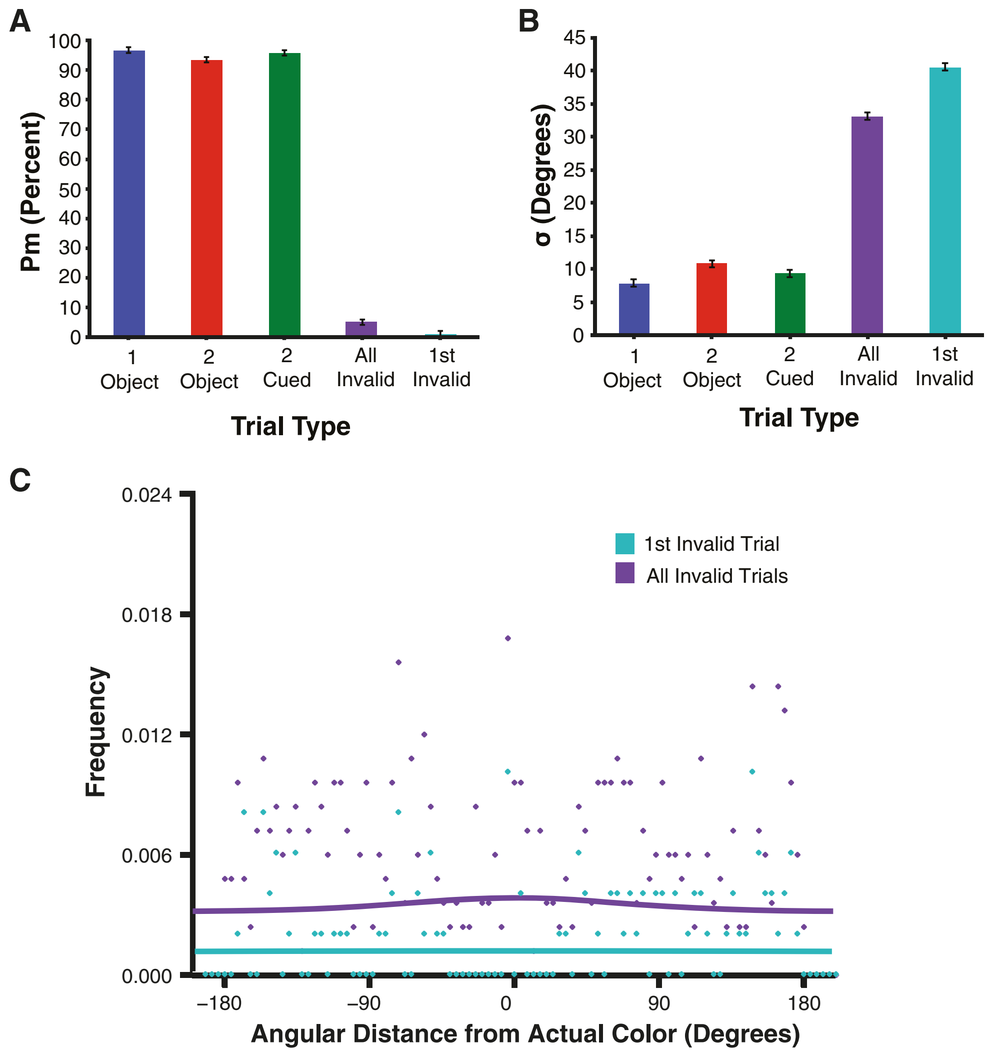

Fig. 3 Results from Experiment 2. **a** *Pm* values (± 2 *SEM*s) and **b** *σ* values (± 2 *SEM*s) from the mixture modeling analysis across all trial types and memory probes. **c** Histogram of all responses to invalidly cued trials (in degrees in color space from the actual color of the probed item), separated into responses to the first trial with an invalid directed-forgetting cue (light blue) and the average of all five invalid trials across all 100 participants (purple). The lines show the mixture model fits

As is shown in the light blue line in Fig. 3c, the distribution of the responses on the first invalid trial yielded an essentially flat line, indicating that participants guessed at the color of the item that they had been cued to forget during the retention interval. The modeling analysis confirmed that the probability that the to-be forgotten representation was still in memory was approximately 1 % ($Pm = .01$). We also ran the modeling analysis on the combination of all five invalid trials (see Fig. 3c, purple line). Again, the mixture model values indicated that representations of the items that participants had been cued to forget were virtually absent from memory when it was probed ($Pm = .07$). When we compared the two-cued trials (trials with valid directed-forgetting cues) against the invalid directed-forgetting cued trials, we found significant differences in both *Pm*, $t(99) = 121.76$, $p < .001$, and σ, $t(99) = 13.18$, $p < .001$. In addition to these analyses, we also used a bootstrapping technique to estimate the confidence intervals for *Pm* (Fougnie & Alvarez, 2011), in which we resampled all of the available invalid-cue trials 1,000 times. Using this method yielded confidence intervals similar to those from the ANOVA (*Pm*: lower bound = .0523, upper bound = .1029), putting the upper bound of the invalid responses almost an order of magnitude lower than the *Pm* of the two-object trials when either item could be tested.

Discussion

The findings of Experiment 2 replicated those of Experiment 1 in showing improvements in precision and in the probability of storage following the presentation of directed-forgetting cues during the retention interval. Moreover, the trials with invalid cues showed that the participants' memory for the item that they had been cued to forget was highly inaccurate, as would de predicted if people were discarding that representation from visual working memory. The findings suggest that the participants guessed at the correct color when presented with these surprise memory probes. Averaging across all five of the invalid trials and the 100 participants showed that nearly all of the responses were guesses, with the mixture modeling only providing a hint of a bump in the distribution at the remembered location. During debriefing, several participants reported that after noticing the possibility of being tested on the cued item, they tried to continue remembering this color during the remaining trials. This likely contributed to the slight increase in the rate of responding around the presented color. These results provide support for the hypothesis that a directed-forgetting cue results in participants completely discarding the representation from visual working memory, provided that nearly all memory cues are valid (see Williams & Woodman, 2012). Although the present findings are consistent with the idea that the to-be-forgotten information is discarded from visual working memory, it is not possible to completely rule out that this information is simply allowed to decay or die in visual working memory. For example, the decay-or-death hypothesis could account for the findings if the rate of decay or death from visual working memory were similar to that of iconic memory (Sperling, 1960). Future research will be needed with improbable invalid cues to determine whether the rate of information loss from visual working memory in the absence of active maintenace might account for the present findings.

General discussion

Our use of mixture modeling during a directed-forgetting procedure revealed how these cues benefit the remaining representations in visual working memory. In Experiment 1, we found that a directed-forgetting cue resulted in a higher probability of retaining the remaining item, as well as greater precision of that representation as compared to when two items were encoded and held in memory throughout the 3-s retention interval. Experiment 2 demonstrated that a cue not only helps to preserve a representation, but also that it probably results in the representation that the cue has indicated can be forgotten being expelled from visual working memory. This shows that maintenance mechanisms can be focused on task-relevant representations during postcuing procedures because the other information is simply removed from the limited-capacity visual working memory store.

An alternative explanation of the present findings—in which the maintenance of certain memory representations does not change, regardless of cues—is that the cues that we presented during the retention intervals reduced interference from the memory probe displays (Makovski & Jiang, 2007; Makovski, Sussman, & Jiang, 2008; Sligte, Scholte, & Lamme, 2008). This interpretation appears unlikely to account for the body of evidence for three reasons. First, the recall probe used in the present experiments was a single outlined box in all conditions. The minimal interference caused by this probe would be similar in all conditions and would not predict the complete absence of information when the participants were probed to report the item that they had been cued to forget. Second, in an unpublished pilot experiment, we found that cues presented earlier in the retention interval—and thus, further from the memory probes—showed greater benefits than did cues presented later in the retention interval. This finding is to be expected if the cues allow a participant to protect certain representations in memory by focusing maintenance mechanisms on them (Griffin & Nobre, 2003; Matsukura, Luck, & Vecera, 2007). In contrast, it is not clear how an account in which the cues prevent interference from memory probes could account for this pattern. Third, we have previously found that cues to forget or remember a subset of the items in visual working memory change the nature of the event-related potential components measured during memory retention intervals (Williams & Woodman, 2012). The observation of these effects a second before any memory test array makes the retrieval-interference explanation untenable.

The findings of the present study have several implications. First, models of memory that explain serial position effects often model the loss of information from working memory as being due to displacement by new information (Raaijmakers & Shiffrin, 1981). Indeed, this argument has an intuitive appeal, as the contents of working memory might often be displaced by our encountering new objects as we move through our environment. However, the present results demonstrate that this is not the only way that information can leave temporary memory. We have shown that humans can expel information from visual working memory in a top-down manner, and that this then benefits the remaining information in visual working memory. This expulsion also has benefits for processing, given that studies have suggested that the maintenance of information in visual working memory might trigger an attentional bias to similar items in the environment (Carlisle & Woodman, 2011a; Olivers, Meijer, & Theeuwes, 2006; Soto, Heinke, Humphreys, & Blanco, 2005), and that these memory representations can even warp our perception of low-level features such as the motion of stimuli (Kang, Hong, Blake, & Woodman, 2011). Though it appears possible to override such attentional biases with high-level goals (e.g., Carlisle & Woodman, 2011b; Woodman & Luck, 2007), if we want to avoid unwanted deployments of attention to items in our visual field that are not currently task-relevant, it would be most effective to keep the representations of task-irrelevant objects out of visual working memory altogether.

The next implication of the present findings is for the physiological economy of the brain. It is clear that task-relevant information is maintained in visual working memory by the sustained firing of neurons (Chelazzi, Miller, Duncan, & Desimone, 1993; Rainer, Asaad, & Miller, 1998), and potentially by the correlated firing between these neurons (Vogel, Woodman, & Luck, 2001). Given this, it would be metabolically economical to discard information from visual working memory, if possible, instead of having displacement remove representations from being actively maintained. The former mechanism would allow the brain to conserve neural activity, whereas the latter would require the brain to fire action potentials constantly, without regard to metabolic costs. The behavioral findings that we have shown here explain previous electrophysiological results (Williams & Woodman, 2012) that have been consistent with the idea that cues presented during retention intervals are used to discard information from visual working memory and to conserve the neural maintenance of information for the most relevant representations.

Author Note This research was supported by the National Eye Institute of the National Institutes of Health (Grant No. RO1-EY019882) and by the National Science Foundation (Grant No. BCS 09-57072 and a graduate research fellowship to M.W.). We thank Kristen Quinn and Kingsley Njoku, who helped with data collection. We also thank Julianna Ianni and Adrienne DiRaddo for their useful feedback and suggestions on this project.

References

Anderson, M. C., & Green, C. (2001). Suppressing unwanted memories by executive control. *Nature, 410,* 366–369.

Anderson, M. C., & Levy, B. J. (2009). Suppressing unwanted memories. *Current Directions in Psychological Science, 18,* 189–194.

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

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision, 10,* 433–436. doi:10.1163/156856897X00357

Carlisle, N. B., & Woodman, G. F. (2011a). Automatic and strategic effects in the guidance of attention by working memory representations. *Acta Psychologica, 137,* 217–225. doi:10.1016/j.actpsy.2010.06.012

Carlisle, N. B., & Woodman, G. F. (2011b). When memory is not enough: Electrophysiological evidence for goal-dependent use of working memory representations in guiding visual attention. *Journal of Cognitive Neuroscience, 23,* 2650–2664.

Chelazzi, L., Miller, E. K., Duncan, J., & Desimone, R. (1993). A neural basis for visual search in inferior temporal cortex. *Nature, 363,* 345–347. doi:10.1038/363345a0

Fougnie, D., & Alvarez, G. A. (2011). Object features fail independently in visual working memory: Evidence for a probabilistic feature-store model. *Journal of Vision, 11*(12), 1–12. doi:10.1167/11.12.3

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

Johnson, H. M. (1994). Processes of successful intentional forgetting. *Psychological Bulletin, 116,* 274–292.

Kang, M.-K., Hong, S. W., Blake, R., & Woodman, G. F. (2011). Visual working memory contaminates perception. *Psychonomic Bulletin & Review, 18,* 860–869.

MacLeod, C. M. (1999). The item and list methods of directed forgetting: Test differences and the role of demand characteristics. *Psychonomic Bulletin & Review, 6,* 123–129. doi:10.3758/BF03210819

Makovski, T., & Jiang, Y. V. (2007). Distributing versus focusing attention in visual short-term memory. *Psychonomic Bulletin & Review, 14,* 1072–1078.

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,* 369–380. 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? *Perception & Psychophysics, 69,* 1422–1434. doi:10.3758/BF03192957

Olivers, C. N. L., Meijer, F., & Theeuwes, J. (2006). Feature-based memory-driven attentional capture: Visual working memory content affects visual attention. *Journal of Experimental Psychology: Human Perception and Performance, 32,* 1243–1265. doi:10.1037/0096-1523.32.5.1243

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision, 10,* 437–442. doi:10.1163/156856897X00366

Raaijmakers, J. G. W., & Shiffrin, R. M. (1981). Search of associative memory. *Psychological Review, 88,* 93–134. doi:10.1037/0033-295X.88.2.93

Rainer, G., Asaad, W. F., & Miller, E. K. (1998). Selective representation of relevant information by neurons in the primate prefrontal cortex. *Nature, 393,* 577–579.

Sligte, I. G., Scholte, H. S., & Lamme, V. A. F. (2008). Are there multiple visual short-term memory stores? *PLos ONE, 3,* e1699. doi:10.1371/journal.pone.0001699

Soto, D., Heinke, D., Humphreys, G. W., & Blanco, M. J. (2005). Early, involuntary top-down guidance of attention from working memory. *Journal of Experimental Psychology: Human Perception and Performance, 31,* 248–261. doi:10.1037/0096-1523.31.2.248

Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs, 74*(11, Whole No. 498), 1–29.

Vogel, E. K., Woodman, G. F., & Luck, S. J. (2001). Storage of features, conjunctions, and objects in visual working memory. *Journal of Experimental Psychology: Human Perception and Performance, 27,* 92–114. doi:10.1037/0096-1523.27.1.92

Williams, M., & Woodman, G. F. (2012). Directed forgetting and directed remembering in visual working memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 38,* 1206–1220.

Woodman, G. F., & Luck, S. J. (2007). Do the contents of visual working memory automatically influence attentional selection during visual search? *Journal of Experimental Psychology: Human Perception and Performance, 33,* 363–377. doi:10.1037/0096-1523.33.2.363

Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representations in visual working memory. *Nature, 453,* 233–235. doi:10.1038/nature06860

Zhang, W., & Luck, S. J. (2009). Sudden death and gradual decay in visual working memory. *Psychological Science, 20,* 423–428. doi:10.1111/j.1467-9280.2009.02322.x