This manuscript is now published in the Journal of Experimental Psychology: General

(DOI: 10.1037/xge0001400)

Abstract

Maintaining perceptual experiences in visual working memory (VWM) allows us to flexibly accomplish various tasks, but some tasks come at a price. For example, comparing VWM representations to novel perceptual inputs can induce inadvertent memory distortions. If these distortions can persist, they may explain why everyday memories often become unreliable after people perform perceptual comparisons (e.g., eyewitness testimony). Here, we conducted two experiments to assess the consequences of perceptual comparisons using real-world objects that were temporarily maintained in VWM ($n = 32$) or recalled from visual long-term memory back into VWM ($n = 30$). In each experiment, young adults reported systematic memory distortions following perceptual comparisons. These distortions increased in magnitude with the delay between encoding and comparisons and were preserved when memories were retrieved again a day later. These findings suggest that perceptual comparisons play a mechanistic role in everyday memory distortions, including situations where memory accuracy is vital.

Keywords

Memory distortion; Perceptual comparisons; Long-term memory; Real-world objects

Public Significance Statement

This study demonstrates that erroneous memory biases formed by comparing one's memory of a visual object to new objects in the surrounding environment may be permanent. In particular, memory biases reported immediately after comparing memories to new percepts were found to be nearly identical in magnitude when the same memories were recalled again 24 hours later, even when individuals expressed high confidence in the accuracy of both reports. The act of perceptual comparison may therefore provide a mechanistic explanation for perseverant errors described in eyewitness lineups where individuals confidently sustain false recognition judgments across time.

Comparing Visual Memories to Similar Visual Inputs Risks Lasting Memory Distortion

Humans can actively maintain a finite amount of visual information in mind after it is no longer available to the senses (Luck & Vogel, 1997; Vogel & Machizawa, 2004). This ability, commonly referred to as visual working memory (VWM), allows us to use past visual experiences—both recent and remote—to accomplish current tasks (Cowan, 2008). For example, suppose you go to the store to buy your partner the jacket that she showed you recently for her birthday. When you arrive at the store, you can recall the appearance of the jacket from visual long-term memory (VLTm) back into VWM and compare this mental representation to the other jackets on display in the store. In order to find the desired jacket—and avoid an unpleasant birthday mishap—you must ensure that your memory of the jacket’s appearance remains intact across comparisons with other similar-looking jackets, as well as across time, as you move from one store to the next.

However, preserving an accurate VWM representation during perceptual comparisons may not be possible. Recent research has demonstrated that comparing VWM representations to novel visual inputs can result in inadvertent distortions to the memories. Across multiple experiments, Fukuda et al. (2022) found that individuals’ VWM reports were biased towards novel visual inputs that were compared with the VWM during a preceding perceptual comparison. These report biases were shown to be larger when individuals endorsed the novel input as similar to their VWM representation than when the novel input was endorsed as dissimilar instead. Using computational modeling, the researchers showed that these retroactive similarity-induced memory biases (SIMB) were likely explained by representational integration between the memory and similar percept, drawing a comparison to serial dependence, in which

prior percepts integrate with current percepts to produce analogous *proactive* biases in perception (e.g., Fischer & Whitney, 2014; Kiyonaga et al., 2017; Manassi & Whitney, 2022).

In a follow-up study, Saito et al. (2022) investigated whether SIMB observed following perceptual comparisons is fundamentally distinct from other forms of retroactive memory bias that have been observed under other task demands. The researchers directly compared SIMB following perceptual comparisons to the biases observed when individuals perceive, but ignore novel inputs during VWM maintenance (e.g., Sun et al., 2017; Teng & Kravitz, 2019; Rademaker et al., 2015) and those observed when individuals maintain multiple memoranda in VWM (e.g., Chunharas et al., 2022; Scotti et al., 2021). Attraction biases were observed in all three tasks but were largest following perceptual comparisons, even after accounting for trial-wise differences in physical stimulus similarity and task differences in memory precision. Critically, the authors found that larger SIMB following perceptual comparisons was the result of bias amplification that occurred in response to perceived similarity, but not perceived dissimilarity, in the probe. Together, these findings have led to the assertion that the task demands associated with performing perceptual comparisons play a causal role in modulating memory biases.

Considering the ubiquity of perceptual comparisons in everyday life and their potential for inducing undesired memory distortion, it is imperative that researchers directly examine the mechanistic reach of perceptual comparisons under more realistic conditions. To this point, extant evidence of SIMB is restricted to simple visual features (e.g., colors, shapes) and has not yet been investigated in complex stimuli that individuals encounter in daily life (e.g., real-world objects). While some studies have shown that memories of real-world objects can be biased away from one another within VWM (e.g., Scotti et al., 2021; Chunharas et al., 2022) and

VLTM (e.g., Chanales et al., 2021; Zhao et al., 2021) to reduce interference, no study has demonstrated that memory representations of real-world objects are capable of being integrated with perceptual representations following perceptual comparisons. On the contrary, some have suggested that the VWM system is capable of representing real-world objects at higher loads and better levels of precision than simple visual stimuli (Brady et al., 2016; but see also Quirk et al., 2020), which may imply their increased resistance to SIMB. These uncertainties surrounding the stimulus-specific nature of SIMB mandate a direct investigation.

Relatedly, no study has directly investigated the durability of SIMB across time. Longstanding models of learning and memory agree that the exchange of information between VWM and VLTm plays a meaningful role in determining the contents of a memory that persist. From a cognitive perspective, perceptual information encoded into VLTm must first pass through the VWM “bottleneck” during initial learning (Atkinson & Shiffrin, 1968; Forsberg et al., 2020; Fukuda & Vogel, 2019). At the time of retrieval, information recalled from VLTm returns to VWM, where it is temporarily maintained for usage in everyday behaviors, such as perceptual comparisons (Cowan, 2008; Fukuda & Woodman, 2017; Sutterer et al., 2019). Neuroscientific models formalize this VWM “buffer” between perception and VLTm as a unique state of the representation, rather than a unique memory system, *per se*. During learning, information is represented “online” in a vulnerable state until the representation is stabilized through the process of consolidation (McGaugh, 2000; Kandel et al., 2014; Dudai, 2004). At retrieval, consolidated representations can return to this labile “online” state again to allow for updating by new information that can elaborate on the memory and promote further learning (Lee, 2009; Hardt et al., 2010). Once updated, the representation returns to its latent consolidated state again, either through reconsolidation (Nader & Hardt, 2009; Tronson & Taylor, 2007) or

the consolidation of a novel trace that contains the updated contents (Winocur & Moscovitch, 2011; Moscovitch, 2007). Across all of these models, VWM plays a fundamental role in the initial formation and continuous elaboration of a memory across time, including in situations where updating occurs undesirably and leads to inaccurate representations (see Schacter et al., 2011; Hardt et al., 2010 for reviews). As such, it is reasonable to expect that VLTm representations are prone to SIMB when retrieved and that SIMB can persist across time when retrieved representations return to their latent, consolidated state. Such a conclusion would provide a novel mechanistic explanation for the pervasive unreliability of memory in critical scenarios, such as eyewitness testimony (Wixted & Wells, 2017).

In the present study, we conducted two experiments showing that perceptual comparisons induce SIMB in real-world objects that are temporarily maintained in VWM (Experiment 1) and those that are retrieved from VLTm (Experiment 2). When individuals retrieved biased memory representations again 24 hours after perceptual comparisons, the magnitude of SIMB was nearly identical to when it was measured immediately after perceptual comparisons. Thus, these results provide direct evidence that deliberate perceptual comparisons can induce lasting distortions in memories of stimuli that individuals encounter in everyday life.

Experiment 1

Before testing the long-term effects of perceptual comparisons, we first needed to establish that comparisons can induce SIMB in complex stimuli that are meaningful and can be plausibly remembered over longer durations. To do this, we adopted the same procedure employed by Fukuda et al. (2022) but used colored real-world objects instead of simple visual features, like colors and shapes.

Method

Transparency and Openness

We report how we determined our sample sizes, all data exclusions, all manipulations, and all measures in the study. All data and analysis code are available at the Open Science Framework (<https://osf.io/z3rgt/>). All analyses were conducted using MATLAB, version R2020a (Mathworks, 2020) and the Psychophysics Toolbox extension, version 3.0.16 (Kleiner et al., 2007). This study's design and analyses were not pre-registered.

Participants

Based on previous work (Fukuda et al., 2022; Saito et al., 2022), we anticipated medium-to-large effect sizes (*Cohen's d* > 0.7). A power analysis with an alpha level of 0.05 and statistical power of 0.95 indicated that we would need at least 29 subjects to obtain such an effect (Faul et al., 2007). Participants were recruited on a weekly basis until the targeted number of subjects was reached. Data were collected from 35 undergraduate students in accordance with the procedures approved by the Research Ethics Board at the University of Toronto. Each participant reported their gender identity and age by hand on the consent form. For gender identity, participants selected between male, female, and prefer not to answer. Race and ethnicity information were not collected. All volunteers had normal or corrected-to-normal visual acuity and did not report color-blindness. Three participants' data were removed because of not following task instructions (one), technical malfunction during the experiment (one), and providing an insufficient number of high-confidence responses ($\leq 10\%$, one), resulting in 32 participants' data being subjected to analysis (25 female, 7 male, mean age = 18.3 years old).

Stimuli and Apparatus

Object stimuli were sampled from a pre-existing database (Brady et al., 2013) and color-rotated in CIE $L^*a^*b^*$ space centered at $a^* = 20$ and $b^* = 38$ with a radius of 60. L^* was set to 70. During the experiment, participants were seated approximately 60cm from an LCD monitor (refresh rate = 60Hz) where colored objects and an object-color wheel were presented on a white background (80 cd/m^2) using the Psychophysics Toolbox in MATLAB (Brainard, 1997).

Procedure

Participants performed six blocks of 30 pseudorandomized trials (**Figure 1C**). Each trial began with a target object ($5.3^\circ \times 5.3^\circ$) presented at the center of the screen for 1500 ms, which participants were instructed to remember as precisely as possible. Target object colors were randomly sampled from the circular color space (see *Stimuli and Apparatus*). Target object presentation was followed by a 5500-ms maintenance interval. At the completion of the maintenance interval, an object-color wheel was presented (**Figure 1A**). The object-color wheel (15.8° diameter) was comprised of 12 equidistant colored exemplars of the target object, whose colors were 30° apart in the circular color space. Using the mouse, participants reported the color of the original target object by clicking on the object-color wheel. Participants were told that they could click between two adjacent exemplars to allow for more precise reporting. After selecting, a response probe ($5.3^\circ \times 5.3^\circ$) was displayed at the center of the screen in the selected color. Participants were able to use the left and right arrow keys to fine-tune the color of the response probe to match what they remembered as precisely as possible. This two-stage continuous report was implemented to minimize interference by the response probe, while preserving participants' ability to report their memory as precisely as possible. Afterwards, they indicated their confidence in the accuracy of their final memory report by pressing one of three

Figure 1. *Experiment 1 Schematic.*

(A) Circular object-color wheel. (B) Probe sampling procedure for each experimental condition. (C) The trial procedure for each condition in the experiment.

keyboard buttons (*high confidence, low confidence, no memory*). The accuracy of the memory report was emphasized and was therefore reported without an imposed time limit.

In two-thirds of the trials, participants completed two perceptual comparisons during the maintenance interval. 500 ms after the offset of the target object, two copies of the target object ($5.3^\circ \times 5.3^\circ$ each) were presented simultaneously as probes on either side of the screen (5.3° from center). Because the perceived similarity of a novel visual input is shown to vary across

trials (Fukuda et al., 2022; Saito et al., 2022), we controlled perceived similarity during each perceptual comparison by rendering one of the probes to be physically similar and the other to be physically dissimilar. The color of the similar probe was randomly sampled $\pm 16 - 45^\circ$ away from the target object. To determine the color of the dissimilar probe, we rotated the sampling window 180° ($\pm 196 - 225^\circ$) and randomly sampled again. Participants were instructed to indicate which of the two probes was most similar to the original target object by pressing the left or right arrow key on the keyboard. The probes remained on the screen for 2000 ms regardless of the report and were followed by a 500-ms blank delay before participants completed a second perceptual comparison on a second pair of object probes. After another 500-ms blank delay following the second perceptual comparison, participants reported the target object color following the same two-step procedure described above. In the *one-way bias condition*, the colors of the similar probes were sampled from the same side of the circular feature space relative to the target object (**Figure 1B**). In the *two-way bias condition*, the colors of the similar probes were sampled from opposite sides of the circular color space relative to the target object (**Figure 1B**). The offset direction of the first similar probe was randomly determined on a trial-by-trial basis.

Analyses

To characterize SIMB in each condition, we computed the response offset for each trial by subtracting the feature value of the final memory report from that of the original item. We aligned the direction of response offsets across trials such that positive values indicated final memory reports that were offset towards the first similar probe (*signed response offsets*). The direction of the response offsets in the baseline condition were randomly assigned. To quantify the size of SIMB across trials, we computed the *bias magnitude* using the mean signed response offset for each condition.

To better isolate the variability in participants' memory reports that was systematically tied to the similar probes, we filtered the data to only include trials where the participant correctly selected the similar probes during both perceptual comparisons ($M = 96.0\%$ of trials in the experimental condition). This helped to reduce the number of trials where participants failed to successfully encode the memory item (e.g., due to a lapse in attention) and were forced to guess during the perceptual comparison. While there may be some theoretical utility in evaluating these guess trials, that question is separate from our current research question and is severely underpowered. Second, from this filtered set of trials, we selected those with high-confidence (HC) final memory reports that were unlikely to be contaminated by guessing ($M = 78.7\%$ baseline trials, $M = 71.3\%$ experimental trials). Notably, when we performed the same analyses without filtering for confidence, the same pattern of results persisted (see *Supplementary Materials*).

We report both frequentist (i.e., t values) and Bayesian (i.e., Bayes Factors) statistics in order to provide evidence in favor of predicted null differences between conditions. BF_{01} indicates evidence in favor of the null hypothesis and BF_{10} indicates evidence in favor of the alternative hypothesis.

Results & Discussion

We began by assessing the effect of the similar probes on the remembered target. As demonstrated previously, perceived similarity in a given probe should result in attraction biases in the report of the target (Fukuda et al., 2022; Saito et al., 2022). As can be seen in **Figure 2A**, signed response distributions in the one-way condition were shifted in the direction of the similar probes, confirming the presence of an attractive bias. When we measured the magnitude of this bias, we found that the systematic attraction was reliable (**Figure 2B**, $M = 3.78^\circ$, 95% CI [1.94,

Figure 2. *Experiment 1 Results.*

(A) Signed response distributions for the baseline and experimental conditions when constrained to trials with high-confidence memory reports. For demonstration purposes, we plotted the proportion of responses for a given signed offset value by calculating the mean response proportion across a 30° window centered at the offset value. Positive offsets indicate memory bias towards the first similar probe. The inset shows a close-up of the peak of each distribution. Shaded regions surrounding the distribution curve indicate within-subject standard errors of the mean (Cousineau, 2005). The vertical dashed line indicates the location of the target centered across trials. (B) Boxplots of the mean signed response error (i.e., bias) in each experimental condition. Positive values indicate memory bias towards the first similar probe. The horizontal line inside each boxplot indicates the median bias across participants. Colored dots to the right of each boxplot indicate the bias for a given participant with corresponding density distributions shown on the left-hand panel of the figure.

* $p < 0.05$, *** $p < 0.001$

5.62°], $t(31) = 4.19$, $p < 0.001$, *Cohen's d* = 0.74, $BF_{10} = 1.29 \times 10^2$). To address the possibility that participants simply biased their reports of the target towards the similar probes on the color wheel to optimize their performance, we conducted a control experiment (Experiment S1) where

we replaced the continuous report at the end of the trial with a two-alternative forced-choice (2AFC) instead. Consistent with the present findings, we found that performing a perceptual comparison increased the likelihood that participants would select an incorrect foil of the target that was sampled in the direction of the similar probe during the 2AFC (see *Supplementary Materials* for a full report). Thus, much like simple visual features, real-world objects are susceptible to systematic biases following perceptual comparisons, even when the individual is highly-confident in the accuracy of their memory report or is re-presented the actual target again at the time of the report.

Next, we assessed the unique contribution of each perceptual comparison to the report bias. Were participants' reports of the target biased by both similar probes during a given trial or were they only biased by the most recent probe? If the target was biased by both similar probes, we should expect that an attraction bias induced by the first similar probe can be counteracted by an opposing attraction bias induced by the second similar probe. As shown in **Figure 2A**, signed response distributions in the two-way condition were not shifted in the direction of the first similar probe and, if anything, showed a small bias towards the second similar probe. When we measured the magnitude of this bias, we found that attraction towards the second similar probe was significant (**Figure 2B**, $M = -1.06^\circ$, 95% CI $[-1.97, -0.14^\circ]$, $t(31) = -2.36$, $p = 0.025$, *Cohen's d* = -0.42, $BF_{10} = 2.06$). If this bias was driven exclusively by the second similar probe with no influence by the first similar probe, we should expect that the absolute magnitudes of the biases in the one- and two-way conditions were comparable. However, we found instead that the absolute magnitude of the bias was larger in the one-way condition than the two-way condition (**Figure 2B**, $\Delta M = 2.72^\circ$, 95% CI $[0.75, 4.70^\circ]$, $t(31) = 2.81$, $p = 0.009$, *Cohen's d* = 0.50, $BF_{10} = 5.06$). This suggests that the small bias towards the second similar probe in the two-way

condition reflected the net outcome of two opposing biases that were induced across the comparisons and that the slightly larger effect of the second similar probe was attributable to differences in the quality of the target representation following a longer maintenance delay and a preceding interference event (i.e., the first comparison) that made the target more vulnerable to interference. Taken together, these findings suggest that real-world objects are susceptible to systematic bias following perceptual comparisons and that the bias induced by a given comparison can serve to correct or exacerbate existing biases.

Experiment 2

In Experiment 2, we measured SIMB in real-world objects retrieved from VLTm. If perceptual comparisons can explain distortions observed in everyday memory use, we should again observe reliable SIMB in these retrieved representations. Moreover, if these anticipated biases reflect durable changes to the representation that are reconsolidated into VLTm, we should expect that SIMB induced by perceptual comparisons will be reported again during a subsequent retrieval episode.

Method

Participants

Based on Experiment 1, we anticipated medium-to-large effect sizes (*Cohen's d* = 0.7). A power analysis with an alpha level of 0.05 and statistical power of 0.95 indicated that we would need at least 29 subjects to obtain such an effect (Faul et al., 2007). Participants were recruited on a weekly basis until the targeted number of subjects was reached. Data were collected from 36 undergraduate students using the same informed consent and screening procedure as Experiment 1. Six participants' data were removed because of not following task instructions

(one), failing to return the second day of the experiment (two), and providing an insufficient number of confident responses ($\leq 10\%$, two) or completed perceptual comparisons ($< 80\%$, one), resulting in 30 participants' data being subjected to analysis (19 female, 11 male, mean age = 18.5 years old).

Stimuli and Apparatus

Identical to Experiment 1.

Procedure

Figure 3A shows a timeline of the two-day procedure used in Experiment 2.

Encoding Task. Participants performed six blocks of an explicit encoding task (**Figure 3B**). In each block, participants were sequentially presented 240 pictures of real-world objects and instructed to remember each object and its color for a subsequent memory test. The same 240 objects were presented in each block in a randomized order, such that each object was presented six times by the end of the task. In every trial, an object ($5.3^\circ \times 5.3^\circ$) was presented at the center of the screen for 1250 ms, followed by a 500-ms blank delay before the next object was presented. To encourage attentiveness on each trial, participants indicated whether the present object had already been shown previously by pressing one of three buttons on the keyboard while the given object was present onscreen (*definitely repeated, maybe repeated, first time presented*).

Memory Test. Participants were tested on their memory of the encoded objects on the same day as encoding (*Test 1*) and 24 hours later (*Test 2*). Each memory test was identical, except for the objects that were tested. As shown in **Figure 3A**, half of the encoded objects were randomly selected to be recalled for the first time at Test 1 (*Immediate Set*) and the remaining half were

Figure 3. *Experiment 2 Schematic.*

(A) The two-day experimental procedure. The times of day are included as an illustrative example and were not identical for every participant in the actual experiment. (B) The object encoding task. (C) The trial procedure for each condition of the memory tests.

recalled for the first time at Test 2 (*Delayed Set*). At each test, half of the objects being recalled for the first time were randomly selected to be tested in the *bias condition*, while the remaining half were tested in the *baseline condition*. This manipulation was performed at each test to assess

whether perceptual comparisons performed in the bias condition reliably distorted representations retrieved from VLTm. At each memory test, 25% of trials included novel objects that were not presented during the encoding task. We included fewer novel objects than studied objects to try and minimize fatigue during the memory tests that could be caused by doubling the total number of objects. Including fewer novel objects than studied objects during each memory test likely induced a response bias favoring “old” responses. However, because this response bias affected all conditions equally, it does not preclude our ability to interpret relative differences in recognition between the Immediate and Delayed Set objects. All of the Immediate Set objects were recalled again at Test 2 in the baseline condition to determine if SIMB at Test 1 persisted across the 24-hour delay.

Participants performed four and eight blocks of 40 pseudorandomized trials at Test 1 and Test 2, respectively (**Figure 3C**). Each trial began with a grayscale object ($5.3^\circ \times 5.3^\circ$) presented at the center of the screen and participants completed an untimed *initial recognition report* where they indicated whether they remembered the object and its color by pressing one of three keyboard buttons (*remember object and color*, *remember object only*, *no memory of object*). In every trial, this report was then followed by a 2750-ms maintenance interval before participants completed a final memory report following the same two-step procedure outlined in Experiment 1. In the bias condition, participants completed a perceptual comparison during the maintenance interval by indicating which of two probe objects ($5.3^\circ \times 5.3^\circ$ each) was most similar to the encoded object by pressing the left or right arrow key while the probes were onscreen. Unlike Experiment 1, the dissimilar probe in Experiment 2 was determined by sampling the color that was 180° opposite the similar probe on each comparison.

Analyses

As in Experiment 1, we again included only trials in which participants correctly selected the similar probe during the perceptual comparison and were highly confident in their final memory report. For the Immediate Set, this made up more than 69% and 63% of the total trials in the baseline and bias conditions at Test 1, respectively. It also led to 64% and 56% of all Immediate Set objects being selected when they were tested again in the baseline condition at Test 2. For the Delayed Set, this made up more than 39% and 33% of the total trials in the baseline and bias conditions at Test 2, respectively. The comparatively lower proportion of analyzed trials in Experiment 2 is not surprising given the much longer retention interval as compared to Experiment 1. Imposing these same strict criteria could, thus, reduce our statistical power. However, we reasoned that eliminating guessing was more important for evaluating the present hypotheses than getting a precise estimate of bias effect size. For one, random guessing can also reduce statistical power by introducing responses that are highly variable and mask the systematic effects of perceptual comparisons. Moreover, guessing itself can be systematic. When individuals fail to retrieve object details, they may try to strategically guess the object's color based on the similar probe. For these reasons, we maintained the same conservative criteria here but report analyses which include all trials in the *Supplementary Materials*.

We computed the bias magnitude in the same manner as Experiment 1. To quantify memory precision before and after the 24-hour delay, we computed the inverse standard deviation of the raw response offsets in the baseline condition for HC trials only, since these were unlikely to be contaminated by guessing. To determine if perceptual comparisons at Test 1 changed participants' subjective estimates of memory quality at Test 2, we compared the proportion of final memory reports made with HC at each memory test and the proportion of

trials where participants indicated that they remembered an encoded object's color during initial recognition reports.

Results & Discussion

SIMB Induced Before or After the Delay

We began by assessing whether object representations recalled for the first time from VLTm at Test 1 (Immediate Set) and Test 2 (Delayed Set) were biased by perceptual comparisons. Signed response distributions are shown in **Figure 4A**. When we measured the magnitude of the bias in each object set, we confirmed that the systematic shifts observed at Test 1 ($M = 5.06^\circ$, 95% CI [2.41, 7.71°], $t(29) = 3.90$, $p < 0.001$, *Cohen's d* = 0.71, $BF_{10} = 5.90 \times 10$) and Test 2 ($M = 10.09^\circ$, 95% CI [6.59, 13.59°], $t(29) = 5.90$, $p < 0.001$, *Cohen's d* = 1.08, $BF_{10} = 8.88 \times 10^3$) both reflected reliable attraction biases towards the similar probe (**Figure 4B**). Interestingly, the bias induced in Delayed Set objects after the overnight delay was significantly larger than the bias induced in Immediate Set objects on the same day as encoding ($\Delta M = 5.03^\circ$, 95% CI [1.31, 8.75°], $t(29) = 2.77$, $p = 0.010$, *Cohen's d* = 0.51, $BF_{10} = 4.60$). Consistent with the representational integration account of SIMB proposed by Fukuda et al. (2022), we found that the larger memory biases in Delayed Set objects compared to Immediate Set objects coincided with significantly lower memory precision in the former (**Figure 4C**, $t(29) = 3.16$, $p = 0.004$, *Cohen's d* = 0.58, $BF_{10} = 1.06 \times 10$; see also Lively et al., 2021). These findings confirm that perceptual comparisons can bias VLTm representations and reveal that more consolidated long-term memories are more vulnerable to distortion than recently-formed ones.

Figure 4. *SIMB Present in Long-Term Memories.*

(A) Signed response distributions for each set of encoded objects (i.e., immediate, delayed) at each memory test (i.e., Test 1, Test 2) when constrained to trials with high-confidence memory reports. For demonstration purposes, we plotted the proportion of responses for a given signed offset value by calculating the mean response proportion across a 30° window centered at the offset value. Positive offsets indicate memory bias towards the similar probe. The inset shows a close-up of the peak of each distribution. Shaded regions surrounding the distribution curve indicate within-subject standard errors of the mean (Cousineau, 2005). The vertical dashed line indicates the location of the target centered across trials. (B) Boxplots of the mean signed response error (i.e., bias) in each set of encoded objects. Positive values indicate memory bias towards the similar probe. (C) Boxplots of the inverse

standard deviation of response errors (i.e., precision) in each set of encoded objects that were recalled in the baseline condition. Larger values indicate better precision. The horizontal line inside every boxplot indicates the median value across participants. Colored dots to the right of each boxplot indicate the value for a given participant with corresponding density distributions shown on the left-hand panel of the figure. T1 = Memory Test 1, T2 = Memory Test 2, SD = standard deviation, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

SIMB Persistence Across the Delay

Before examining the durability of SIMB across time from Test 1 to Test 2, we first assessed whether participants' metacognitive estimates of their memory at Test 2 were influenced by previously retrieving the memory at Test 1. To do this, we computed the proportion of Immediate Set Objects recalled in the baseline conditions at Test 1 and Test 2 where participants reported that they remembered the object and its color during the initial recognition judgment and reported the color of the object with high confidence during the memory report. As can be seen in **Figure 5A**, these metacognitive estimates decreased slightly for Immediate Set objects from Test 1 to Test 2, suggesting that some objects became less precise or were forgotten between tests (initial recognition: ($\Delta M = 0.13$, 95% CI [0.07, 0.20], $t(29) = 4.09$, $p < 0.001$, *Cohen's d* = 0.75, $BF_{10} = 9.34 \times 10$; confidence: $\Delta M = 0.04$, 95% CI [0.00, 0.09], $t(29) = 1.88$, $p = 0.071$, *Cohen's d* = 0.34, $BF_{01} = 1.10$). However, compared to Delayed Set objects being recalled for the first time at Test 2, participants reported that they remembered Immediate Set objects and their color more often at Test 2 (**Figure 5A**; $\Delta M = 0.20$, 95% CI [0.15, 0.24], $t(29) = 8.74$, $p < 0.001$, *Cohen's d* = 1.60, $BF_{10} = 9.02 \times 10^6$) and with greater confidence ($\Delta M = 0.25$, 95% CI = [0.22, 0.29], $t(29) = 13.67$, $p < 0.001$, *Cohen's d* = 2.49, $BF_{10} = 1.89 \times 10^{11}$). These patterns suggest that Immediate Set objects were more

Figure 5. *SIMB Persistence Across 24-hour Delay.*

(A) The proportion of baseline trials for Immediate and Delayed Set Objects at Test 1 and Test 2 where participants reported that they remembered the object and its color during the initial recognition judgment (*left*) and the proportion of trials where participants reported the color of the object with high confidence (*right*). Height of the bars indicate mean values across participants with errors bars indicating within-subject standard errors of the mean (Cousineau, 2005). Small colored circles represent the mean value for a given participant. (B) Signed response distributions at Memory Test 2 for Immediate Set Objects that were recalled with high confidence at both memory tests. For demonstration purposes, we plotted the proportion of responses for a given

signed offset value by calculating the mean response proportion across a 30° window centered at the offset value. Positive offsets indicate memory bias towards the similar probe that was presented at Memory Test 1. The inset shows a close-up of the peak of each distribution. Shaded regions surrounding the distribution curve indicate within-subject standard errors of the mean (Cousineau, 2005). The vertical dashed line indicates the location of the target centered across trials. (C) Boxplots of the mean signed response error (i.e., bias) at Memory Test 1 and 2 for Immediate Set Objects that were recalled with high confidence at both memory tests. Positive values indicate memory bias towards the similar probe that was presented at Memory Test 1. The horizontal line inside each boxplot indicates the median bias across participants. Colored dots to the right of each boxplot indicate the bias for a given participant with corresponding density distributions shown on the left-hand panel of the figure. T1 = Memory Test 1, T2 = Memory Test 2, ns = not significant, ** $p < 0.01$, *** $p < 0.001$

accessible to participants during Test 2 because of their prior retrieval at Test 1, consistent with the mnemonic benefits typically conferred by retrieval practice (Roediger & Karpicke, 2006; Roediger & Butler, 2011).

From here, we moved to assess whether Immediate Set objects had recovered from the biases that were induced by the similar probe at Test 1. To account for objects that were not successfully accessed during one of the tests or were forgotten between tests, we compared the magnitude of the bias between Tests 1 and 2 for Immediate Set objects that were reported with high confidence at both tests. We found a nearly identical bias at each test (**Figure 5C**; Test 1: $M = 4.84^\circ$; Test 2: $M = 4.78^\circ$; $\Delta M = 0.06^\circ$, 95% CI [-2.02, 2.14°], $t(29) = 0.06$, $p = 0.953$, *Cohen's* $d = 0.01$, $BF_{01} = 5.14$), suggesting that the bias induced during Test 1 in Immediate Set objects was encoded into VLTm and accessed again at Test 2. Together, these results show that, far from transient, SIMB induced by perceptual comparisons is robustly stored in VLTm, remaining every bit as powerful a day after the comparison as compared to seconds afterwards.

General Discussion

In previous experiments, using VWM representations of simple visual stimuli in perceptual comparisons causally distorted the memory representations (Fukuda et al., 2022; Saito et al., 2022). Yet, no study has definitively established whether these distortions reflect lasting changes to memory content. In line with prominent theories and empirical evidence, we hypothesized that systematic distortions caused by perceptual comparisons would occur when VLTm representations are recalled back “online” into VWM (Cowan, 2008; Fukuda & Woodman, 2017; Sutterer et al., 2019; Lee, 2009; Hardt et al., 2010; Schacter et al., 2011) and that these distortions would persist through the process of reconsolidation and be reported again during subsequent retrieval (Atkinson & Shiffrin, 1968; Forsberg et al., 2020; Fukuda & Vogel, 2019; Nader & Hardt, 2009; Tronson & Taylor, 2007; Winocur & Moscovitch, 2011; Moscovitch, 2007).

First, we established that real-world objects temporarily maintained in VWM are susceptible to SIMB, despite potential differences in how well the VWM system is able to represent these realistic items compared to simple visual stimuli (Brady et al., 2016; Quirk et al., 2020). We then tested the generalizability of SIMB to VLTm by asking participants to encode real-world objects into VLTm in anticipation of immediate and delayed (i.e., 24 hours) memory tests that included perceptual comparisons. We found that VLTm representations were distorted by comparisons at both tests, but more so at the delayed test, when memory precision was lower, consistent with the basic assumptions of representational integration (Fukuda et al., 2022; Saito et al., 2022). Importantly, we also found that objects reported during the immediate test showed SIMB of comparable magnitude again at the delayed test, despite the fact that individuals were confident in the accuracy of both reports. Together, these findings suggest that using VLTm

representations in perceptual comparisons risks lasting memory distortions that arise outside of metacognitive awareness.

Lasting Biases Following Perceptual Comparisons Imply a True Cognitive Mechanism

In nearly every study of VWM bias, including SIMB, the nature of the bias is characterized using a psychophysical approach where participants report the contents of their memory by selecting from a continuous wheel (e.g., Sun et al., 2017; Teng & Kravitz, 2019; Rademaker et al., 2015; Chunharas et al., 2022; Scotti et al., 2021). However, the use of continuous estimation reports to measure memory bias welcomes concerns about trivial response strategies that are capable of producing biases in the report *independent* of the memory representation (e.g., Chunharas et al., 2022). For example, in the present paradigm, participants could have fine-tuned their reports of the target towards the features of the similar probe—either intentionally or unintentionally—in attempt to optimize their performance or communicate their understanding of probe similarity. This fine-tuning strategy would produce an observable bias in participants’ behavioral reports without requiring any change in the underlying memory representation.

However, as alluded to in Experiment 1, such an account cannot provide a straightforward explanation for the results of our control experiment (Experiment S1), where we replaced the continuous estimation report at the end of the trial with a two-alternative forced-choice (2AFC) report instead where participants selected between the correct target and a biased foil of the target (see *Supplementary Materials* for full report). We found that participants were more likely to select an incorrect foil of the target that was sampled *towards* the similar probe if participants performed a perceptual comparison beforehand. This confirmed that biased reports following perceptual comparisons are not specific to continuous estimation tasks where

participants can leverage information about the probe and the color wheel to optimize their performance. Critically, we also found that performing a perceptual comparison did not necessarily make participants more likely to select the correct target during the 2AFC when the foil was sampled *away* from the similar probe. This lack of symmetry in response bias between the foil sampling conditions eliminates the possibility of an analogous report strategy during the 2AFC task where participants universally selected the option that looked more like the similar probe. Instead, we speculate that performing a perceptual comparison before the 2AFC did not increase the likelihood of selecting the correct target amongst a “repulsed” foil because the attractive memory bias induced by the comparison always rendered both options in the 2AFC different from one’s memory. Together, the findings between Experiments 1 and S1 cannot be readily explained by trivial report strategies and, if anything, show that performing perceptual comparisons can cause individuals to mistakenly endorse an incorrect stimulus, even when the correct stimulus is present (c.f., eyewitness testimony, Wixted & Wells, 2017).

While the present results cannot be clearly explained by strategic biasing of the report, they may be explained by a bona fide bias that occurred at the time of the report. Studies have shown that visual perception is serially-dependent, such that previously-viewed stimuli can bias one’s current perception (e.g., Fischer & Whitney, 2014; Kiyonaga et al., 2017; Manassi & Whitney, 2022). Is SIMB actually the result of a *proactive* bias in the response probe that was shown during the continuous report rather than a retroactive bias in the memory representation? Findings in the present study clearly undermine this possibility. If participants’ perception of the response probe was being attracted towards the similar probe with no bias to the target memory, then participants’ reports of the preserved target would show a reliable repulsion bias, as participants would need to correct for the attractive serial dependence in their perception of the

response probe in order to match the probe to their memory of the target. While a repulsive serial dependence in the response probe would be more consistent with the present findings, serial dependencies in color perception are shown to be invariably attractive when the current percept is within 100° of the prior one in the feature space (see Barbosa & Compte, 2020, for a recent review). This leaves retroactive SIMB as the better explanation.

With that acknowledged, studies of serial dependence may still provide useful insight into the mechanisms that are responsible for producing SIMB. For example, in a recent study by Sheehan & Serences (2022), the authors show that serial dependencies in participants' behavioral reports are the result of a "read-out" mechanism that translates *repulsive* neural biases into *attractive* behavioral biases. Given the opposite direction of causality between serial dependence and SIMB, the finding of opposite neural and behavioral biases in serial dependence does not immediately challenge existing evidence that attractive memory biases are caused by integration between neural representations (Fukuda et al., 2022; Saito et al., 2022; see also Rademaker et al., 2015, 2019; Bae et al., 2015 for similar conceptualizations). However, these findings could be used to speculate about the exact time course of SIMB. It may be the case that the integration processes responsible for producing SIMB did not occur until individuals attempted to access the memory representation again during the report that followed perceptual comparisons. If so, representational integration can be characterized as a "read-out" process at the time of the report that is primed by the performance of a perceptual comparison. More importantly, this could also imply that choosing not to report a memory after performing a perceptual comparison may prevent representational integration from taking place, thereby protecting the memory against the biasing effect of the perceptual comparison. Future work should seek to identify when the bias

occurs by implementing neural decoding approaches that are capable of tracking the contents of VWM from the time of the perceptual comparison until the completion of the behavioral report.

A related possibility is that representational integration did, in fact, occur at the time of the comparison, but the process of reporting the memory representation afterwards is what allowed the bias to persist across time. In each of the present experiments, the behavioral report of the target required continuously manipulating a response probe to match memory content (Experiments 1-2) or endorsing a fixed probe as matching memory content (Experiment S1). In each case, the observer was directly exposed to a perceptual representation of the target in a biased state. This exposure to a biased perceptual representation of the target during the report may have caused the bias to persist in VLTm, either by strengthening the biased memory and increasing its likelihood of reconsolidation (Nader & Hardt, 2009; Tronson & Taylor, 2007) or by creating an additional trace of the biased memory that made it more accessible during subsequent retrieval (Winocur & Moscovitch, 2011; Moscovitch, 2007). A straightforward test of this hypothesis would be to measure the persistence of SIMB without an immediate memory report following the comparison or by using an alternative type of report that does not require the observer to directly perceive the biased representation.

Perceptual Comparisons and Real-World Memory Distortion

Based on the present findings, perceptual comparisons may reflect a cognitive mechanism that can account for real-world memory distortions, such as those described in eyewitness testimony, where individuals explicitly compare their memories to novel perceptual inputs (Wixted & Wells, 2017). However, there are several meaningful differences between the present paradigm and everyday behavior. We highlight some of these differences below to establish the mechanistic reach of the current evidence and to motivate future research.

In the present experiments, we relied on a psychophysical approach to collect fine-grained estimates of memory accuracy and precision for a single visual feature that has been well-validated (i.e., color, Zhang & Luck, 2008). In doing so, we show that systematic memory biases can be induced even when individuals retain object-based representations in order to meet subsequent retrieval demands. An immediate next step will be to assess SIMB when multiple task-relevant features in the novel input are manipulated. Face stimuli would be ideal candidates for meeting this criterion and for providing a direct link to eyewitness testimony (for preliminary evidence, see Plummer et al., 2021).

Relatedly, memory distortions here were operationalized as biases along a continuous task-relevant feature. One may argue that the relatively small magnitude of these changes is unlikely to significantly disrupt everyday functioning. This may be true in some cases. However, as shown in Experiment 1, biases arising from perceptual comparisons can be additive when the same memory is used repeatedly. This is consistent with related work showing that false identifications in eyewitness testimony are more frequent following repeated recognition decisions (Stebly & Dysart, 2016; Wixted et al., 2021).

Nonetheless, many large-scale memory distortions in the real-world are thought to arise from source misattribution (Gershman et al., 2013; Carpenter & Schacter, 2011). This may be analogous to “swapping” effects described in VWM literature, where individuals report the features of a non-target item due to probabilistic confusion (Bays et al., 2009) or as an intentional strategy to compensate for poor memory quality (Pratte, 2018). Can perceptual comparisons induce swapping as well? Preliminary evidence suggests that it may depend on the judgment made during the comparison. Here, memories were biased when individuals compared the similarity between memories and inputs. However, when memories and novel inputs are

endorsed as being the *same* rather than *similar*, memories appear to be replaced by (i.e., “swapped with”) novel inputs instead (Saito et al., 2022b). These preliminary findings are broadly consistent with the view that source misattribution and other related memory distortions are the result of memory updating processes that are typically viewed as functionally-adaptive for behavior (Schacter et al., 2011).

With that being said, similarity and sameness are fundamentally distinct. While both require the perception of featural overlap between items, sameness indicates the lack of any (even subtle) detectable differences between items. As such, it may be difficult to induce false endorsements of sameness between VWMs and novel inputs when individuals encode both items with reasonable precision. If so, then swapping observed following endorsements of sameness may be explained by poor memory quality and not by comparisons, *per se*. However, one promising circumvention is to manipulate contextual details that influence comparisons without influencing the stimuli. For example, eyewitness lineups result in more false identifications if the suspect resembles the perpetrator description more than the fillers in that lineup (Fitzgerald et al., 2013; Colloff et al., 2016). In this example, diagnostic criteria in the perceptual environment shift the likelihood of perceiving sameness without changing the physical appearance of the suspect. Once that suspect is endorsed as the perpetrator, individuals tend to confidently sustain this decision, suggesting that the memory of the original perpetrator was replaced by the falsely-identified suspect (Wixted et al., 2015; Roediger et al., 2012). Future work should incorporate contextual manipulations into the present paradigm to elucidate the long-term consequences of endorsing sameness between memories and inputs.

Constraints on Generality

We show that young adults’ memories of real-world objects can be systematically biased

following perceptual comparisons with similar visual inputs and that these biases can persist across time in LTM. From this, we posit that perceptual comparisons may serve as a viable mechanism for explaining perseverant memory errors described in applied fields of study, such as eyewitness lineups, where comparisons are performed explicitly. However, it is unknown how emotional factors present in eyewitness contexts might change how SIMB manifests outside of a controlled laboratory setting. In a recent review by Glomb (2022), the author highlights the lack of a theoretical framework that can account for seemingly-opposed patterns of memory performance demonstrated by witnesses of emotionally-charged events. Thus, in order to draw stronger inferences about the generalizability of SIMB to eyewitness scenarios, laboratory studies of SIMB should include manipulations to emotion that closely mimic those shown to produce paradoxical patterns of memory performance in the real world.

References

- Atkinson, R. C., & Shiffrin, R. M. (1968). Human memory: A proposed system and its control processes. In K. W. Spence & J. T. Spence (Eds.), *The Psychology of Learning and Motivation: Advances in Research and Theory* (Vol. 2, pp. 89–195). New York: Academic Press.
- Bae, G.-Y., Olkkonen, M., Allred, S. R., & Flombaum, J. I. (2015). Why some colors appear more memorable than others: A model combining categories and particulars in color working memory. *Journal of Experimental Psychology: General*, 144(4), 744–763.
<https://doi.org/10.1037/xge0000076>
- Barbosa, J., & Compte, A. (2020). Build-up of serial dependence in color working memory. *Scientific Reports*, 10(1), 1-7. <https://doi.org/10.1038/s41598-020-67861-2>
- 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(10), 7-7.
<https://doi.org/10.1167/9.10.7>
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, 10(4), 433–436. <https://doi.org/10.1163/156856897X00357>
- Brady, T. F., Konkle, T., Gill, J., Oliva, A., & Alvarez, G. A. (2013). Visual long-term memory has the same limit on fidelity as visual working memory. *Psychological Science*, 24(6), 981-990. <https://doi.org/10.1177/0956797612465439>
- Brady, T. F., Störmer, V. S., & Alvarez, G. A. (2016). Working memory is not fixed-capacity: More active storage capacity for real-world objects than for simple stimuli. *PNAS*

- Proceedings of the National Academy of Sciences of the United States of America*, 113(27), 7459–7464. <https://doi.org/10.1073/pnas.1520027113>
- Carpenter, A. C., & Schacter, D. L. (2017). Flexible retrieval: When true inferences produce false memories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 43(3), 335. <https://dx.doi.org/10.1037%2Fxl0000340>
- Chanales, A. J., Tremblay-McGaw, A. G., Drascher, M. L., & Kuhl, B. A. (2021). Adaptive repulsion of long-term memory representations is triggered by event similarity. *Psychological Science*, 32(5), 705-720. <https://doi.org/10.1177/0956797620972490>
- Chunharas, C., Rademaker, R. L., Brady, T. F., & Serences, J. T. (2022). An adaptive perspective on visual working memory distortions. *Journal of Experimental Psychology: General*, 151(10), 2300–2323. <https://doi.org/10.1037/xge0001191>
- Colloff, M. F., Wade, K. A., & Strange, D. (2016). Unfair lineups make witnesses more likely to confuse innocent and guilty suspects. *Psychological Science*, 27(9), 1227-1239. <https://doi.org/10.1177%2F0956797616655789>
- Cousineau, D. (2005). Confidence intervals in within-subject designs: A simpler solution to Loftus and Masson's method. *Tutorials in quantitative methods for psychology*, 1(1), 42-45. <https://doi.org/10.20982/tqmp.01.1.p042>
- Cowan, N. (2008). What are the differences between long-term, short-term, and working memory? In W. Sossin, J.-C. Lacaille, V. F. Castellucci, & S. Belleville (Eds.), *Progress in brain research: The essence of memory* (Vol. 169, pp. 323–338). Amsterdam: Elsevier/Academic Press.

- Dudai, Y. (2004). The neurobiology of consolidations, or, how stable is the engram. *Annual Review of Psychology*, 55, 51–86.
<https://doi.org/10.1146/annurev.psych.55.090902.142050>
- Faul, F., Erdfelder, E., Lang, A. G., & Buchner, A. (2007). G* Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, 39(2), 175–191. <https://doi.org/10.3758/BF03193146>
- Fischer, J., & Whitney, D. (2014). Serial dependence in visual perception. *Nature Neuroscience*, 17(5), 738–743. <https://doi.org/10.1038/nn.3689>
- Fitzgerald, R. J., Price, H. L., Oriet, C., & Charman, S. D. (2013). The effect of suspect-filler similarity on eyewitness identification decisions: A meta-analysis. *Psychology, Public Policy, and Law*, 19(2), 151. <https://psycnet.apa.org/doi/10.1037/a0030618>
- Forsberg, A., Guitard, D., & Cowan, N. (2020). Working memory limits severely constrain long-term retention. *Psychonomic Bulletin & Review*, 1–11. <https://doi.org/10.3758/s13423-020-01847-z>
- Fukuda, K., Pereira, A. E., Saito, J. M., Tang, T. Y., Tsubomi, H., & Bae, G. Y. (2022). Working memory content is distorted by its use in perceptual comparisons. *Psychological Science*, 33(5), 816–829. <https://doi.org/10.1177/09567976211055375>
- Fukuda, K., & Vogel, E. K. (2019). Visual short-term memory capacity predicts the “bandwidth” of visual long-term memory encoding. *Memory & Cognition*, 47(8), 1481–1497.
<https://doi.org/10.3758/s13421-019-00954-0>

- Fukuda, K., & Woodman, G. F. (2017). Visual working memory buffers information retrieved from visual long-term memory. *Proceedings of the National Academy of Sciences*, 114(20), 5306-5311. <https://doi.org/10.1073/pnas.1617874114>
- Gershman, S. J., Schapiro, A. C., Hupbach, A., & Norman, K. A. (2013). Neural context reinstatement predicts memory misattribution. *Journal of Neuroscience*, 33(20), 8590-8595. <https://doi.org/10.1523/JNEUROSCI.0096-13.2013>
- Glomb, K. (2022). How to improve eyewitness testimony research: Theoretical and methodological concerns about experiments on the impact of emotions on memory performance. *Psychological Research*, 86(1), 1-11. <https://doi.org/10.1007/s00426-021-01488-4>
- Hardt, O., Einarsson, E. Ö., & Nader, K. (2010). A bridge over troubled water: Reconsolidation as a link between cognitive and neuroscientific memory research traditions. *Annual Review of Psychology*, 61, 141–167. <https://doi.org/10.1146/annurev.psych.093008.100455>
- Lee, J. L. C. (2009). Reconsolidation: Maintaining memory relevance. *Trends in Neurosciences*, 32(8), 413–420. <https://doi.org/10.1016/j.tins.2009.05.002>
- Lively, Z., Robinson, M. M., & Benjamin, A. S. (2021). Memory Fidelity Reveals Qualitative Changes in Interactions Between Items in Visual Working Memory. *Psychological Science*, 32(9), 1426-1441. <https://doi.org/10.1177/0956797621997367>
- Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, 390(6657), 279-281. <https://doi.org/10.1038/36846>

- Kandel, E. R., Dudai, Y., & Mayford, M. R. (2014). The molecular and systems biology of memory. *Cell*, 157(1), 163-186. <https://doi.org/10.1016/j.cell.2014.03.001>
- Kiyonaga, A., Scimeca, J. M., Bliss, D. P., & Whitney, D. (2017). Serial dependence across perception, attention, and memory. *Trends in Cognitive Sciences*, 21(7), 493-497. <https://doi.org/10.1016/j.tics.2017.04.011>
- 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(2), 369. <https://psycnet.apa.org/doi/10.1037/0278-7393.34.2.369>
- Makovski, T., & Pertzov, Y. (2015). Attention and memory protection: Interactions between retrospective attention cueing and interference. *Quarterly Journal of Experimental Psychology*, 68, 1735–1743. <https://doi.org/10.1080/17470218.2015.1049623>
- Manassi, M., & Whitney, D. (2022). Illusion of visual stability through active perceptual serial dependence. *Science Advances*, 8(2), eabk2480. <https://doi.org/10.1126/sciadv.abk2480>
- McGaugh, J. L. (2000). Memory: A century of consolidation. *Science*, 287(5451), 248–251. <https://doi.org/10.1126/science.287.5451.248>
- Moscovitch M. (2007). Memory: Why the engram is elusive. In Roediger H. L., Dudai Y., Fitzpatrick S. M. (Eds.), *Science of memory: Concepts* (pp. 17–21). New York, NY: Oxford University Press.
- Nader, K., & Hardt, O. (2009). A single standard for memory: The case for reconsolidation. *Nature Reviews Neuroscience*, 10(3), 224-234. <https://doi.org/10.1038/nrn2590>

Plummer, M., Hellerstedt, R., Gibson, S., Simons, J., & Bergstrom, Z. M. (2021, October 25).

Active Recognition Attempts Induce Updating of Face Memories. *PsyArXiv*.

<https://doi.org/10.31234/osf.io/63qnj>

Pratte, M. S. (2019). Swap errors in spatial working memory are guesses. *Psychonomic Bulletin & Review*, 26(3), 958-966. <https://doi.org/10.3758/s13423-018-1524-8>

Quirk, C., Adam, K. C., & Vogel, E. K. (2020). No evidence for an object working memory capacity benefit with extended viewing time. *Eneuro*, 7(5).

<https://doi.org/10.1523%2FENEURO.0150-20.2020>

Rademaker, R. L., Bloem, I. M., De Weerd, P., & Sack, A. T. (2015). The impact of interference on short-term memory for visual orientation. *Journal of Experimental Psychology: Human Perception and Performance*, 41(6), 1650–1665.

<https://doi.org/10.1037/xhp0000110>

Rademaker, R. L., Chunharas, C., & Serences, J. T. (2019). Coexisting representations of sensory and mnemonic information in human visual cortex. *Nature Neuroscience*, 22(8), 1336–1344. <https://doi.org/10.1038/s41593-019-0428-x>

Roediger III, H. L., & Butler, A. C. (2011). The critical role of retrieval practice in long-term retention. *Trends in Cognitive Sciences*, 15(1), 20-27.

<https://doi.org/10.1016/j.tics.2010.09.003>

Roediger, H. L. III, & Karpicke, J. D. (2006). Test-enhanced learning: Taking memory tests improves long-term retention. *Psychological Science*, 17(3), 249–255.

<https://doi.org/10.1111/j.1467-9280.2006.01693.x>

- Roediger, H. L., Wixted, J. T., & DeSoto, K. A. (2012). The curious complexity between confidence and accuracy in reports from memory. In L. Nadel & W. Sinnott-Armstrong (Eds.), *Memory and law* (pp. 84–117). New York, NY: Oxford University Press.
- Saito, J. M., Duncan, K., & Fukuda, K. (2023, January 28). Comparing Visual Memories to Similar Visual Inputs Risks Lasting Memory Distortion. Retrieved from osf.io/z3rgt
- Saito, J. M., Kolisnyk, M., & Fukuda, K. (2022). Perceptual comparisons modulate memory biases induced by new visual inputs. *Psychonomic Bulletin & Review*. Advance online publication. <https://doi.org/10.3758/s13423-022-02133-w>
- Saito, J. M., Bae, G., & Fukuda, K. (2022, September 15). Judgments During Perceptual Comparisons Predict Distinct Forms of Memory Updating. *PsyArXiv*. <https://doi.org/10.31234/osf.io/pfx6>
- Schacter, D. L., Guerin, S. A., & St. Jacques, P. L. (2011). Memory distortion: An adaptive perspective. *Trends in Cognitive Sciences*, 15(10), 467–474. <https://doi.org/10.1016/j.tics.2011.08.004>
- Scotti, P. S., Hong, Y., Leber, A. B., & Golomb, J. D. (2021). Visual working memory items drift apart due to active, not passive, maintenance. *Journal of Experimental Psychology: General*, 150(12), 2506–2524. <https://doi.org/10.1037/xge0000890>
- Sheehan, T. C., & Serences, J. T. (2022). Attractive serial dependence overcomes repulsive neuronal adaptation. *PLoS Biology*, 20(9), e3001711. <https://doi.org/10.1371/journal.pbio.3001711>

- Stebly, N. K., & Dysart, J. E. (2016). Repeated eyewitness identification procedures with the same suspect. *Journal of Applied Research in Memory and Cognition*, 5(3), 284-289.
<https://doi.org/10.1016/j.jarmac.2016.06.010>
- Sun, S. Z., Fidalgo, C., Barense, M. D., Lee, A. C. H., Cant, J. S., & Ferber, S. (2017). Erasing and blurring memories: The differential impact of interference on separate aspects of forgetting. *Journal of Experimental Psychology: General*, 146(11), 1606–1630.
<https://doi.org/10.1037/xge0000359>
- Sutterer, D. W., Foster, J. J., Serences, J. T., Vogel, E. K., & Awh, E. (2019). Alpha-band oscillations track the retrieval of precise spatial representations from long-term memory. *Journal of Neurophysiology*, 122(2), 539-551.
<https://doi.org/10.1152/jn.00268.2019>
- Teng, C., & Kravitz, D. J. (2019). Visual working memory directly alters perception. *Nature Human Behaviour*, 3(8), 827-836. <https://doi.org/10.1038/s41562-019-0640-4>
- Tronson, N. C., & Taylor, J. R. (2007). Molecular mechanisms of memory reconsolidation. *Nature Reviews Neuroscience*, 8(4), 262-275. <https://doi.org/10.1038/nrn2090>
- Vogel, E. K., & Machizawa, M. G. (2004). Neural activity predicts individual differences in visual working memory capacity. *Nature*, 428(6984), 748-751.
<https://doi.org/10.1038/nature02447>
- Wang, B., Theeuwes, J., & Olivers, C. N. (2018). When shorter delays lead to worse memories: Task disruption makes visual working memory temporarily vulnerable to test interference. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 44(5), 722. <https://psycnet.apa.org/doi/10.1037/xlm0000468>

Winocur, G., & Moscovitch, M. (2011). Memory transformation and systems consolidation.

Journal of the International Neuropsychological Society, 17(5), 766–780.

<https://doi.org/10.1017/S1355617711000683>

Wixted, J. T., Mickes, L., Clark, S. E., Gronlund, S. D., & Roediger III, H. L. (2015). Initial eyewitness confidence reliably predicts eyewitness identification accuracy. *American Psychologist*, 70(6), 515. <https://psycnet.apa.org/doi/10.1037/a0039510>

Psychologist, 70(6), 515. <https://psycnet.apa.org/doi/10.1037/a0039510>

Wixted, J. T., Wells, G. L., Loftus, E. F., & Garrett, B. L. (2021). Test a witness's memory of a suspect only once. *Psychological Science in the Public Interest*, 22, 1S-18S.

<https://doi.org/10.1177%2F15291006211026259>

Wixted, J. T., & Wells, G. L. (2017). The relationship between eyewitness confidence and identification accuracy: A new synthesis. *Psychological Science in the Public Interest*, 18(1), 10-65. <https://doi.org/10.1177%2F1529100616686966>

Interest, 18(1), 10-65. <https://doi.org/10.1177%2F1529100616686966>

Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representations in visual working memory. *Nature*, 453(7192), 233-235. <https://doi.org/10.1038/nature0686>

memory. *Nature*, 453(7192), 233-235. <https://doi.org/10.1038/nature0686>

Zhao, Y., Chanale, A. J., & Kuhl, B. A. (2021). Adaptive memory distortions are predicted by feature representations in parietal cortex. *Journal of Neuroscience*, 41(13), 3014-3024.

<https://doi.org/10.1523/JNEUROSCI.2875-20.2021>