Working memory content is distorted by its use in perceptual comparisons

Keisuke Fukuda^{1,2}, April E. Pereira³, Joseph M. Saito¹, Ty Y. Tang⁴, Hiroyuki Tsubomi⁵, & Gi-Yeul Bae⁴

¹University of Toronto, ²University of Toronto Mississauga, ³University of Waterloo, ⁴Arizona State University, & ⁵University of Toyama

Corresponding author: Keisuke Fukuda

Email: keisuke.fukuda@utoronto.ca

Keywords

Working memory; Memory distortion; Individual differences

Author Contributions

K.F., A.E.P., and H.T. designed the study, A.E.P. and J.M.S. collected data, K.F. analyzed data,

G.B. conducted the modeling analysis, all the authors wrote the manuscript.

Abstract

Visual information around us is rarely static. To carry out a task in such a dynamic environment, we often have to compare current visual input with our working memory representation of the immediate past. However, little is known about what happens to a working memory (WM) representation when it is compared with perceptual input. Here, we tested university students and found that perceptual comparisons retroactively bias working memory representations toward subjectively-similar perceptual inputs. Furthermore, using computational modeling and individual differences analyses, we found that representational integration between WM representations and perceptually-similar input underlies this similarity-induced memory bias. Together, our findings highlight a novel source of WM distortion and suggest a general mechanism that determines how WM representations interact with new perceptual input.

Statement of Relevance

Imagine a person who witnessed a hit-and-run traffic accident. When the witness tries to identify the license plate number of the car involved, a bus occludes his sight of the car. When the bus passes, he needs to find the car by comparing the cars he sees now with the memory of the car that committed the accident. In doing so, he might assume that his memory of the car remains intact as it is compared with the cars he sees now. However, the present study shows that this assumption is simply not valid: As a memory representation of the immediate past (i.e., working memory representation) is compared with new perceptual information, the memory representation becomes systematically biased towards perceptually-similar input. Thus, our findings reveal an overlooked vulnerability in working memory representations and suggest that the witness may misreport the license plate of a similar-but-different car to the police.

Introduction

Many of our daily activities rely on our ability to maintain accurate mental representations of the immediate past to guide our behavior. For example, consider a good Samaritan who witnessed a phone falling out of a passerby's pocket on a busy street. The good Samaritan would look down to pick up the phone in hopes of returning it to its owner. After doing so, he would need to find the owner again by comparing his memory of the owner's appearance to the other people he sees on the street. Past studies have demonstrated that working memory (WM) can actively maintain a small amount of task-relevant information (e.g., the owner's appearance) over a brief delay when the information is not perceptually present (e.g., looking down to pick up the phone; Luck & Vogel, 2013; Ma et al., 2014). However, little is known as to whether a WM representation remains intact as it is used for memory-guided perceptual comparisons.

Though one might assume that a WM representation remains intact as it is used, recent studies show that the memory representation interacts with new perceptual inputs that are physically similar to it (e.g., Sun et al., 2017; Teng & Kravitz, 2019). For example, a previous study asked participants to maintain a color in WM while completing an RSVP task in which a stream of letters was presented overlaid on a differing color patch (Teng & Kravitz, 2019). The researchers found that the color representation being maintained in WM was distorted by the task-irrelevant color patch especially when the two colors were physically-similar to each other. This result demonstrates that the perception of task-irrelevant distractor can distort a WM representation.

However, important questions remain regarding the interactions between WM representations and perceptual input. First, considering that the RSVP task drew attention away from the WM representation, it is possible that the distortion occurred because the WM representation was not protected from the memory-irrelevant distractor or because the maintenance of the memory item was not prioritized (Matsukura et al., 2007). It is also not clear whether the distortion was driven by the replacement of the memory representation with a new perceptual input (Bays et al., 2009) or it was a bona fide distortion of the memory representation. Thus, it remains to be seen whether and how a WM representation would be distorted when the memory representation remains in the focus of attention, and is directly compared with a new task-relevant stimulus.

Second, given a positive correlation between the physical and perceived similarity of stimuli, it is not clear whether the memory distortion is a purely stimulus-driven phenomenon that is determined by a physical similarity of the new stimulus, or it is modulated by participants' subjective perception of the new stimulus.

To test these questions, we had participants remember a single visual feature (i.e., color or shape) in WM and used this information to make a comparison with new stimuli. In three experiments, we found that a WM representation was distorted, especially when participants perceived the new stimuli as similar to the memory representation. Critically, we found that observed memory biases were larger following perceived similarity than perceived dissimilarity, even after controlling for the physical similarity of the new stimulus. We further validated the systematic effect of subjective similarity by demonstrating that memory biases induced by an initial similarity judgment were able to be "cancelled out" by an additional judgment that biased the memory representation in an opposing way. Computational modeling of these similarity-induced memory biases

determined that the observed changes in WM representations were driven by representational integration with the perceptually-similar input. Thus, our findings reveal an overlooked window of vulnerability in WM representations and suggest a general mechanism that determines how a new perceptual input interacts with a WM representation (e.g., Kiyonaga et al., 2017; Rademaker et al., 2019; Sun et al., 2017).

Experiment 1A and 1B

To examine the consequence of comparing a WM representation with perceptual inputs, participants were asked to remember a simple stimulus (i.e., shape in Experiment 1A or color in Experiment 1B) and compared it with a new perceptual input (see Figure 1 for shape task). Subsequently, participants reported their memory of the original stimulus as precisely as possible. Here, we predicted that the WM representation of the original stimulus would be biased towards a new input as a result of the perceptual comparison, especially when the new input was subjectively similar to the memory representation.

Method

Participants

We anticipated a medium-to-large effect size (*Cohen's d* = 0.8) based on related studies. A power analysis determined that at least 15 samples would be necessary for the statistical power of 0.8 (Faul et al., 2007). We recruited 16 (11 females, mean age = 23.2 years old) participants for Experiment 1A and 16 (10 females, mean age = 22.6 years old) participants for Experiment 1B. Participants provided informed consents and the study was approved by the University of Toronto REB.

Apparatus

Stimuli were generated in MATLAB using the Psychophysics Toolbox (Brainard, 1997) and were presented at 60Hz on an LCD monitor. Viewing distance was approximately 60 cm. The same apparatus was used for all reported experiments.

Stimuli

The shape stimulus set used in Experiment 1A contained 360 shapes whose circular continuity was empirically validated (Li et al., 2019). The color stimuli used in Experiment 1B comprised of 360 equally-spaced color values sampled from CIE $L^*a^*b^*$ space centered at $a^* = 20$ and $b^* = 38$ with radius of 60. L^* was set to 70.

Procedure

Participants performed one baseline condition and one perceptual comparison condition (see Figure 1 for the shape task). In the baseline condition, participants were presented with one stimulus (a shape in Experiment 1A: $3.8^\circ \times 3.8^\circ$; a colored circle in Experiment 1B: 5.2° diameter) for 800 ms. After a 3200-ms retention interval, participants were presented with a circular stimulus wheel (15.4° diameter). For Experiment 1A (shape task), the stimulus wheel was composed of 18 equidistant shapes (20° apart in the shape space) that were placed 20° apart on an imaginary circle. For Experiment 1B (color task), it was composed of 360 equidistant colors in the color space. To report the original stimulus using the stimulus wheel, they first selected an item that

Figure 1. The stimulus spaces and the schematic of Experiment 1A

Panel A represents the stimulus space used for Experiment 1A (shape task) and Experiment 1B (color task). Panel B shows the schematic of Experiment 1A.

best matched the original stimulus from the stimulus wheel by rotating a line indicator using the left and right arrow keys on the keyboard. Once the indicator pointed towards the desired item, participants pressed the spacebar to confirm the selection. The selected item then appeared at the center of the screen for optional refinements using the left and right arrow keys. When satisfied, participants indicated their confidence in the accuracy of their report (*1 = high confidence, 2 = low confidence, 3 = no confidence*) by pressing one of three keys on the keyboard. No time limit was imposed on the color report in order to emphasize the accuracy of the report.

In the perceptual comparison condition, one probe was presented 800 ms after the offset of the original stimulus, and participants indicated whether it looked similar or dissimilar to the original stimulus using a key press. The probe stayed on the screen for 1600 ms. The probe was randomly sampled from 16-105° away from the original stimulus. After an 800-ms delay following the probe presentation, participants reported the original memory item as in the baseline condition. Participants performed 12 experimental blocks, and each block contained 12 baseline and 36 experimental trials.

Analysis

For each trial, a signed response offset was calculated as the response offset in the direction of the probe. Thus, a positive value indicated a response offset towards the probe, and a negative value indicated an offset away from it. For the baseline trials, the sign of the response offset was randomly assigned. This created distributions of response

offsets with information about the direction of the bias with respect to the probe. To quantify the magnitude of the bias, we computed the mean of the signed response offsets for each condition. We focused on trials with high-confidence memory reports but the same analyses performed across all trials produced the same pattern of results (see Supplementary Materials).

To isolate the effect of the perceived similarity on the memory bias, we first identified the probe distances that resulted in both similar and dissimilar judgments on separate trials within subjects (*ambivalent probe distances*). For each ambivalent probe distance, we calculated the mean bias magnitude following similar and dissimilar judgments separately. The mean bias magnitudes following similar and dissimilar judgments were then averaged across all ambivalent probe distances. This procedure isolated the effect of perceived probe similarity on the bias magnitude while equating the physical probe similarity (i.e., the distance in the feature space).

Results

Figure 2. Results of Experiment 1A and 1B.

(A) Probability distributions of signed response offsets for high-confidence (HC) reports in Experiment 1A (see E for Experiment 1B). For demonstration purposes, the response proportion for a given signed response offset was computed by calculating the mean response proportion across a 30° window centered around the signed response offset. (B) Violin plots of the mean response bias for HC memory reports in Experiment 1A (see F for Experiment 1B). Thick horizontal lines indicate the mean across participants. Positive values indicate memory bias towards the probe. (C) Probability distributions of signed response offsets for HC reports for ambivalent probe trials in Experiment 1A (see G for Experiment 1B). (D) Violin plots of the mean response bias for HC memory reports for ambivalent probe trials in Experiment 1A (see H for Experiment 1B).

Both shape (Experiment 1A) and color (Experiment 1B) memory reports exhibited an attraction bias ($> 0^\circ$) towards the probe when the probe was judged to be similar to the memory item (shape: $t(15) = 4.70$, $p < 0.001$, *Cohen's d* = 1.18, Figure 2A&B; color: $t(15) = 5.47$, $p < 0.001$, *Cohen's d* = 1.37, Figure 2E&F). We also found some evidence for an attraction bias when the probe was judged to be dissimilar to the original memory item (shape: $t(15) = 1.48$, $p = 0.16$, *Cohen's d* = 0.37; color: $t(15) = 2.93$, $p = 0.01$, *Cohen's d* = 0.73), but the bias magnitude was significantly smaller than when the probe was judged to be similar (similar vs. dissimilar shape: $t(15) = 3.14$, $p <$

0.01, *Cohen's d* = 0.78; similar vs. dissimilar color: $t(15) = 2.98, p < 0.01$, *Cohen's d* = 0.74). The trials with ambivalent probe distances produced larger attraction bias when the probes were perceived to be similar rather than dissimilar to the original memory items (shape task: $t(15) = 3.01, p < 0.01$, *Cohen's d* = 0.75, Figure 2C&D; color task: $t(15) = 4.50, p < 0.01$, *Cohen's d* = 1.13, Figure 2G&H).

Discussion

Experiment 1A and 1B revealed that a WM representation is biased as a consequence of a perceptual comparison with a new input. Particularly, this memory bias was a result of attraction toward a perceptually-similar input. Furthermore, the magnitude of this similarity-induced memory bias was greater when ambivalent probes were judged to be similar to the memory item than dissimilar to it. This suggested that perceived similarity exaggerated the memory bias.

Experiment 2

Experiment 1A and 1B demonstrated that a WM representation is biased by a perceptual comparison with a subjectively-similar stimulus. This suggests that a WM representation may no longer be accurate after it is compared with such an input. Given that many of our daily activities rely on accurate WM representations, we examined whether the biased WM representation can be “corrected” by an additional perceptual comparison. To do so, participants performed two intervening perceptual comparisons with probes sampled from either the same side and opposite sides of the memory item in the stimulus space (Figure 3B). Importantly, in Experiment 1A and 1B, we found that the perceived similarity for a given featural distance was variable across trials. To increase the consistency in the subjective similarity judgments, we presented a pair of probes consisting of a physically-similar probe (i.e., 16-45° away from the target) and a probe that was maximally distant from the physically-similar probe (i.e., 180° away from the similar probe), and had participants choose which probe was more similar to the memory representation.

Method

Participants

Based on the results of Experiment 1A and 1B (*Cohen's d* > 0.73), we anticipated a medium-to-large effect size (*Cohen's d* = 0.73). A power analysis determined that 22 samples were necessary to achieve such an effect size with a statistical power of 0.9 (Faul et al., 2007). After providing written informed consent to protocols approved by the Research Ethics Board (REB) of the University of Toronto, 32 individuals (24 females, mean age = 19.3 years old) participated. The data from four participants were removed due to incompleteness of the experiment (one), failure to follow instructions (one), or too infrequent high-confidence memory reports (i.e., < 15% for the baseline conditions, two). As a result, 28 participants' data were subjected to the analyses.

Procedure

Participants performed four blocks of the color task and four blocks of the shape task in a pseudo-randomized order. In the baseline conditions (Figure 3A), participants

WORKING MEMORY CONTENT DISTORTED BY USE

remembered the original stimulus over a 500-ms (*short delay*) or a 5500-ms (*long delay*) delay interval. After the delay, a circular stimulus wheel was presented, and participants clicked on an item that best matched the original stimulus with a mouse. The selected item then appeared at the center of the screen for optional refinements using the left and right arrow keys. Subsequently, participants reported their confidence in the accuracy of their report in the same manner as previous experiments.

Figure 3. Schematic of Experiment 2.

(A) Schematic of color task in Experiment 2 (B) Two perceptual comparison conditions in the color task.

In the perceptual comparison conditions, participants performed two intervening perceptual comparisons 500-ms after the stimulus offset. In each comparison, two probes were presented on each side of the screen (5.2° from the center of the screen), and participants reported which probe looked more similar to the original stimulus by pressing either the left or the right arrow key. One of the probes was randomly sampled from $\pm 16\text{--}45^\circ$ away from the original stimulus (i.e., similar probe). The other probe was sampled 180° away from the similar probe (i.e., dissimilar probe). The two probes remained on the screen for 2000 ms regardless of the report. After the offset of the first pair of probes, a 500-ms delay followed, and the second pair of probes was presented for another similarity judgment (2000 ms). After another 500-ms delay, participants reported the original stimulus in the same manner as the baseline conditions.

The two perceptual comparison conditions differed in how similar probes were sampled (Figure 3B). In the *same-side probe* condition, similar probes for the first and the second similarity judgments were sampled from the same side of the stimulus space relative to the memory item. In the *opposite-side probe* condition, similar probes for the first and the second perceptual comparisons were sampled from opposite sides of the stimulus space relative to the memory item. Participants completed 40 trials for each condition in a pseudo-randomized order.

Analysis

For each trial, a signed response offset was calculated as the response offset in the direction of the similar probe in the first probe pair. Thus, a positive value indicates a response offset towards the first similar probe, and a negative value indicates an offset away from the first similar probe. For the baseline trials, the sign of the response offset was randomly assigned. The magnitude of the memory bias was quantified as the mean of the signed response offsets for each condition. We focused on trials with high-confidence memory reports, but the same analyses performed across all trials produced the same pattern of results (see Supplementary Materials).

Results

As can be seen from Figure 4A and 4C, the distribution for the *same-side probe* condition was shifted towards the positive direction whereas the distribution for the *opposite-side probe* condition was shifted toward the negative direction. Indeed, the mean signed response offset for the *same-side probe* condition was significantly greater than zero (shape task: $t(27) = 7.17, p < 0.001$, *Cohen's d* = 1.36, Figure 4B; color task: $t(27) = 6.30, p < 0.001$, *Cohen's d* = 1.19, Figure 4D), indicating that the memory representation was attracted toward the similar probes. However, the mean signed response offset for the *opposite-side probe* condition was numerically smaller than zero (shape task: $t(27) = -0.78, p = 0.44$, *Cohen's d* = -0.15, color task: $t(27) = 2.20, p = 0.04$, *Cohen's d* = -0.42), indicating that the memory representation was, if anything, biased away from the first similar probe or biased toward the second similar probe. The magnitude of the bias (i.e., absolute values) for the *same-side probe* condition was statistically greater than the magnitude of the bias for the *opposite-side probe* conditions, (shape task: $t(27) = 7.46, p < 0.001$, *Cohen's d* = 1.41; color task: $t(27) = 3.87, p < 0.001$, *Cohen's d* = 0.73).

Figure 4. Experiment 2 Results.

(A) Probability distribution of signed high-confidence (HC) response offsets for the shape task (see C for the color task). For demonstration purposes, the response proportion for a given signed response offset was computed by calculating the mean response proportion across a 30° window centered around the signed response offset. Positive offsets indicate memory bias toward the first similar probe. The broken lines indicate the within-subject standard errors of the mean. (B) Violin plots of the high-confidence (HC) memory bias for the shape task (see D for the color task). Thick horizontal lines indicate the mean across participants. Positive values indicate memory bias toward the first similar probe.

Discussion

We replicated the similarity-induced memory bias. More importantly, we demonstrated that an additional similarity judgment could counteract the memory bias induced by a previous similarity judgment, thus “correcting” the memory representation.

Modeling of the similarity-induced memory bias

What computational mechanism underlies the similarity-induced memory bias? Here, we tested two alternative models. The joint density model assumes that the WM representation is integrated with the new probe representation if the probe is perceptually similar to the memory representation (Figure 5A, see Dubé et al., 2014 for a similar conceptualization). This integration is accomplished via the formation of a joint probability distribution of the two representations, resulting in a biased memory representation that is shifted towards a similar probe representation. Alternatively, the mixture density model assumes that participants’ reports are sometimes based on the probe representation, and that this confusion is more likely when the probe is perceptually similar to the memory representation (Figure 5B). Thus, the mixture density model predicts that the original memory representation is not biased, but that final reports

can appear to be biased as a result of probabilistic confusion that occurs when a probe is proximal to the memory item.

Method

Detailed descriptions of each model can be found in the Supplementary Material. Here, we provide a summary.

For both the joint density model and the mixture density model, we assume that both the memory (X_M) and the probe (X_P) representations follow von Mises distribution (denoted by ϕ) centered at the stimulus value (S_M, S_P) with some precision (κ_M, κ_P).

$$p(X_M|S_M) = \phi(X_M|S_M, \kappa_M) \quad (1)$$

$$p(X_P|S_P) = \phi(X_P|S_P, \kappa_P) \quad (2)$$

When the probe item (Equation 2) is perceived to be similar to the memory item (Equation 1), the joint density (JD) model integrates the two representations in the following manner:

$$p(X_{JD}|S_M, S_P) = \frac{p(X_M|S_M)p(X_P|S_P)}{\sum p(X_M|S_M)p(X_P|S_P)} \quad (3)$$

Figure 5. Two competing models of similarity-induced memory bias

(A) Joint density model. Noisy representation of a working memory item (X_M) is assumed to follow a von Mises distribution centered at the sample stimulus (S_M , dashed vertical line in black) with some precision (κ_M) (Left). Noisy representation of a probe item is also assumed to follow a von Mises distribution centered at the probe stimulus (S_P , dashed vertical line in blue) with some precision (κ_P) (Center). When observers decide that the memory item and the probe item are perceptually similar, then the two items are integrated to produce a joint distribution (shown in red, Right). As a result, the mean of the joint distribution is biased toward the probe item, as indicated by the dashed red vertical line. (B) Mixture density model. The initial memory representation and the probe representation are assumed to follow von Mises distributions as in the joint density model. However, this model assumes that some proportion of the memory reports are based on the probe representation. This can be accounted for by the mixture parameter (α). The original memory representation in this model is not biased, but the mean of the mixture distribution can be shifted toward the probe item (the dashed red vertical line).

The joint density model (Equation 3) has four parameters. The two parameters for the center of the memory and the probe distributions were set by the actual stimulus values (S_M , S_P). The precision parameter for memory item (κ_M) was obtained by fitting a standard WM model (Zhang & Luck, 2008) to response offsets in the baseline condition in Experiment 2. However, we fit the precision for the probe item within the model (κ_P). Thus, the joint density model has only one free parameter.

We fit the joint density model to each trial and each participant separately for Experiment 2A and 2B data sets. We only used high-confidence memory reports to avoid contamination by guessing or lapses of attention. On a given trial for a given participant, we constructed a joint density distribution using S_M and S_P for that trial and κ_M for the participant and fit the model by finding a probe precision (κ_P) that minimizes the difference between the average human response error collapsed across all the trials (including all the probe distances) and the average simulated response errors across all the simulated responses.

The mixture density model (Equation 4) combines the two distribution via a mixture parameter (α). Namely, this model assumes that the final memory reports are either memory-based reports or probe-based reports. The proportion of each is determined by the mixture parameter.

$$p(X_{\text{Mix}} | S_M, S_P) = \alpha p(X_M | S_M) + (1-\alpha)p(X_P | S_P) \quad (4)$$

All the other aspects of this model were identical to the joint density model except that this model has an additional free parameter (α). α was set to vary between 0 (= 0% memory-based reports) and 1 (= 100% memory-based reports).

Results

Figure 6A shows simulated response offset distributions from the joint density model and the mixture density model along with observed human data (Experiment 2A & 2B). The peak of the simulated response distribution from the joint density model was shifted positively as in the human data. However, the distributions from the mixture density model were positively skewed without the shift. This result suggests that the observed biases in the human data were more likely to be driven by representational integration than probabilistic confusion. Formal model comparisons using sum of log-likelihood, AIC, and BIC measurements unanimously preferred the joint density model (Table 1). To provide additional support for the joint density model, we computed the correlation between the simulated and observed bias magnitudes as a function of the 16 discretized probe distances. The results revealed significant positive correlations for both stimuli, even though the model fit was not done separately for each distance ($r = .79$, $p < 0.001$ for color; $r = .54$, $p = 0.019$ for shape, Figure 6B).

Figure 6. The observed and simulated memory bias in Experiment 1A and 1B (A) Probability distribution of signed human response error, best-fit joint density (JD) model responses, and best-fit mixture density model response for the shape and color tasks in Experiment 1A & 1B. (B) Correlation of bias magnitudes across the 16 probe distances between the simulated responses from the joint density model and the observed responses from human participants in the shape and color tasks in Experiment 1A and 1B. Positive values indicate a bias towards the probe item. Vertical error bars indicate ± 1 bootstrapped standard error of the simulated responses. Horizontal error bars indicate ± 1 bootstrapped standard error of human data.

Data	Model (free parameters)	Σ Log Likelihood	AIC	BIC
Shape (Exp. 1A)	Joint Density (κ_p)	-179.7096	391.4191	482.6021
	Mixture Density (α, κ_p)	-294.6754	653.3508	835.7167
Color (Exp. 1B)	Joint Density (κ_p)	-301.0448	634.0897	723.0092
	Mixture Density (α, κ_p)	-476.1254	1016.251	1194.09

Table 1. Model comparisons. Likelihood values for each model. Boldface indicates the preferred model. κ_p = Precision of probe representation, α = mixture parameter.

Discussion

We found convincing evidence that representational integration underlies similarity-induced memory bias. Although probabilistic confusion might have taken place, our results suggest that such a mechanism was likely not a major contributor to the

observed bias. In Experiment 3, we sought to obtain empirical support for the representational integration account by examining individual differences in the magnitude of the memory bias.

Experiment 3

One novel prediction of the joint density model is that individuals with lower WM precision should exhibit larger memory bias because of greater representational overlap between the memory and probe representations (Figure 5A). To test this, we examined the correlation between individuals' WM precision and the magnitude of the similarity-induced memory bias using a variant of Experiment 1.

Method

Participants

A power analysis based on the effect size obtained in Experiment 1 ($r < -0.32$ between WM precision and the bias magnitude) determined that at least 99 samples would be necessary for the statistical power of 0.9. Following the same informed consent procedure as in previous experiments, 124 (97 females, mean age = 18.6 years old) individuals participated. The data from 14 participants were removed based on the same exclusion criteria used in Experiment 1.

Procedure

The experiment was identical to Experiment 1 except for the following: there was one experimental condition with one intervening similarity judgment and one delay-matched baseline condition. The stimulus onset asynchrony (SOA) between the memory item and the response wheel was set to 4000 ms for both conditions. Participants performed 4 blocks consisting of 15 trials of the baseline and experimental conditions (30 trials per block) in a pseudo-randomized order.

Analysis

Memory precision was estimated by fitting the concentration parameter (κ) of a von Mises distribution to the baseline condition. Again, we focused on trials with high-confidence (HC) response offsets. The precision estimates for participants were then correlated with bias magnitudes obtained in the HC responses of the experimental condition.

Results

We again replicated the attraction bias (see Supplementary Materials). More importantly, individuals with high visual WM precision exhibited a smaller memory bias than those with low precision ($r(107) = -0.37, p < 0.001$ for color; $r(105) = -0.31, p = 0.001$ for shape; Figure 7), as predicted by the joint density model.

Figure 7. Experiment 3 results.

Empty circles indicate individual participants. Crosses indicate participants excluded from the correlation as outliers.

General Discussion

We demonstrated that the use of information retained in WM for a perceptual comparison systematically biased the WM representation. This bias was particularly large when a new visual input was subjectively judged to be similar to the content of visual WM. The joint density model, which attributes the bias to integration between the original memory representation and a subjectively-similar input, successfully captured the bias and ruled out the possibility that the bias was driven by probabilistic confusions between the memory item and the new perceptual input. We also provided empirical support of this representational integration account by demonstrating that individuals with lower visual WM precision exhibited a larger memory bias than those with higher precision.

Past studies have shown that a WM representation can be distorted by multiple sources of task-irrelevant distraction during maintenance (Lorenc et al., 2018; Rademaker et al., 2015; Sun et al., 2017; Teng & Kravitz, 2019). Although our study is critically different from these past studies, in that participants here utilized a maintained WM representation to evaluate new perceptual inputs, some might wonder whether bottom-up visual interference was sufficient to explain the memory bias observed here. Our results suggest that this is unlikely, since the memory bias observed here was strongly modulated by the perceived similarity of the probe even after controlling for its physical similarity to the remembered stimulus. This finding demonstrates a causal role of perceived similarity in modulating the magnitude of memory bias. Our results may instead suggest that the visual interference effect reported in previous studies could have been caused in part by participants incidentally perceiving task-irrelevant items as similar to an original memory item.

Past studies demonstrated that current visual WM representations are biased by past visual experience, a phenomenon known as serial dependence in visual perception (Bae & Luck, 2019, 2020; Cicchini et al., 2018; Czoschke et al., 2019; Fischer et al., 2020; Fischer & Whitney, 2014; Kiyonaga et al., 2017). Although our findings bear

noticeable similarities to serial dependence, they are critically different in the direction of causality. Whereas serial dependence occurs proactively by past experience influencing new visual perception even when the past experience is no longer relevant to the current goal of the task, the bias effect observed in the present study occurred retroactively by new task-relevant visual inputs influencing an existing memory representation. That being said, both phenomena might be explained by our representational integration account. That is, participants in the serial dependence experiment might hold in mind the information in the previous trial while they perceive a new stimulus for the current trial (Bae & Luck, 2019). The co-existence of these representations might prompt participants to perceive similarity between them, which could then lead to representational integration. Future research should investigate this possibility.

Given the pivotal role that WM plays in long-term memory retrieval, our findings of WM distortion propose a plausible mechanism of long-term memory distortion. Mounting evidence suggests that information retrieved from long-term memory can be manipulated by suggestive information provided at the time of retrieval (Frederick et al., 2011; Loftus, 1979; Loftus & Loftus, 1974; Roediger III et al., 1996; Schacter et al., 2011). According to influential memory theories, long-term memories are “brought back” into WM when they are retrieved (Atkinson & Shiffrin, 1968; Cowan, 2001; Fukuda & Woodman, 2017). Our findings suggest that if new information provided at the time of retrieval is deemed similar to a long-term memory retrieved into WM, the retrieved long-term memory might be biased toward the similar information. Future studies should examine this possibility directly by having participants perform similarity judgments between new visual inputs and retrieved long-term memories.

Acknowledgments

We thank Steven J. Luck for the constructive comments on the manuscript.

This research was supported by the Natural Sciences and Engineering Research Council (RGPIN-2017-06866) and the Connaught New Researcher Award.

The data are posted publicly at Open Science Framework (<https://osf.io/5v7sg/>).

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). Academic Press.
- Bae, G. Y., & Luck, S. J. (2019). Reactivation of Previous Experiences in a Working Memory Task. *Psychol Sci*, 30(4), 587-595.
<https://doi.org/10.1177/0956797619830398>
- Bae, G. Y., & Luck, S. J. (2020). Serial dependence in vision: Merely encoding the previous-trial target is not enough. *Psychon Bull Rev*, 27(2), 293-300.
<https://doi.org/10.3758/s13423-019-01678-7>
- Bays, P. M., Catalao, R. F., & Husain, M. (2009). The precision of visual working memory is set by allocation of a shared resource. *J Vis*, 9(10), 7.1-11.
<https://doi.org/10.1167/9.10.7>
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spat Vis*, 10(4), 433-436.
- Cicchini, G. M., Mikellidou, K., & Burr, D. C. (2018). The functional role of serial dependence. *Proc Biol Sci*, 285(1890). <https://doi.org/10.1098/rspb.2018.1722>
- Cowan, N. (2001). The magical number 4 in short-term memory: a reconsideration of mental storage capacity. *Behav Brain Sci*, 24(1), 87-114; discussion 114-185.
- Czoschke, S., Fischer, C., Beitner, J., Kaiser, J., & Bledowski, C. (2019). Two types of serial dependence in visual working memory. *Br J Psychol*, 110(2), 256-267.
<https://doi.org/10.1111/bjop.12349>
- Dubé, C., Zhou, F., Kahana, M. J., & Sekuler, R. (2014). Similarity-based distortion of visual short-term memory is due to perceptual averaging. *Vision Res*, 96, 8-16.
<https://doi.org/10.1016/j.visres.2013.12.016>
- Faul, F., Erdfelder, E., Lang, A., & 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.
- Fischer, C., Czoschke, S., Peters, B., Rahm, B., Kaiser, J., & Bledowski, C. (2020). Context information supports serial dependence of multiple visual objects across memory episodes. *Nat Commun*, 11(1), 1932. <https://doi.org/10.1038/s41467-020-15874-w>
- Fischer, J., & Whitney, D. (2014). Serial dependence in visual perception. *Nat Neurosci*, 17(5), 738-743. <https://doi.org/10.1038/nn.3689>
- Frenda, S. J., Nichols, R. M., & Loftus, E. F. (2011). Current Issues and Advances in Misinformation Research. *Current Directions in Psychological Science*, 20(1), 20-23.
- Fukuda, K., & Woodman, G. F. (2017). Visual working memory buffers information retrieved from visual long-term memory. *Proc Natl Acad Sci U S A*, 114(20), 5306-5311. <https://doi.org/10.1073/pnas.1617874114>
- Kiyonaga, A., Scimeca, J. M., Bliss, D. P., & Whitney, D. (2017). Serial Dependence across Perception, Attention, and Memory. *Trends Cogn Sci*, 21(7), 493-497.
<https://doi.org/10.1016/j.tics.2017.04.011>

- Li, A. Y., Liang, J. C., Lee, A. C. H., & Barense, M. D. (2019). The validated circular shape space: Quantifying the visual similarity of shape. *J Exp Psychol Gen*.
<https://doi.org/10.1037/xge0000693>
- Loftus, E. F. (1979). The malleability of human memory. *Am Sci*, 67(3), 312-320.
- Loftus, G. R., & Loftus, E. F. (1974). The influence of one memory retrieval on a subsequent memory retrieval. *Mem Cognit*, 2(3), 467-471.
<https://doi.org/10.3758/BF03196906>
- Lorenc, E. S., Sreenivasan, K. K., Nee, D. E., Vandenbroucke, A. R. E., & D'Esposito, M. (2018). Flexible Coding of Visual Working Memory Representations during Distraction. *J Neurosci*, 38(23), 5267-5276.
<https://doi.org/10.1523/JNEUROSCI.3061-17.2018>
- Luck, S. J., & Vogel, E. K. (2013). Visual working memory capacity: from psychophysics and neurobiology to individual differences. *Trends Cogn Sci*, 17(8), 391-400.
<https://doi.org/10.1016/j.tics.2013.06.006>
- Ma, W. J., Husain, M., & Bays, P. M. (2014). Changing concepts of working memory. *Nat Neurosci*, 17(3), 347-356. <https://doi.org/10.1038/nn.3655>
- Matsukura, M., Luck, S. J., & Vecera, S. P. (2007). Attention effects during visual short-term memory maintenance: protection or prioritization? *Percept Psychophys*, 69(8), 1422-1434.
- Rademaker, R. L., Bloem, I. M., De Weerd, P., & Sack, A. T. (2015). The impact of interference on short-term memory for visual orientation. *J Exp Psychol Hum Percept Perform*, 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. *Nat Neurosci*, 22(8), 1336-1344. <https://doi.org/10.1038/s41593-019-0428-x>
- Roediger III, H. L., Jacoby, D., & McDermott, K. B. (1996). Misinformation Effects in Recall: Creating False Memories through Repeated Retrieval. *Journal of Memory and Language*, 35, 300-318. <https://doi.org/10.1006/jmla.1996.0017>
- Schacter, D. L., Guerin, S. A., & St Jacques, P. L. (2011). Memory distortion: an adaptive perspective. *Trends Cogn Sci*, 15(10), 467-474.
<https://doi.org/10.1016/j.tics.2011.08.004>
- 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. *J Exp Psychol Gen*, 146(11), 1606-1630.
<https://doi.org/10.1037/xge0000359>
- Teng, C., & Kravitz, D. J. (2019). Visual working memory directly alters perception. *Nat Hum Behav*, 3(8), 827-836. <https://doi.org/10.1038/s41562-019-0640-4>
- Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representations in visual working memory. *Nature*, 453(7192), 233-235.
<https://doi.org/10.1038/nature06860>

Figure captions

Figure 1. The stimulus spaces and the schematic of Experiment 1A
Panel A represents the stimulus space used for Experiment 1A (shape task) and Experiment 1B (color task). Panel B shows the schematic of Experiment 1A.

Figure 2. Results of Experiment 1A and 1B.
(A) Probability distributions of signed response offsets for high-confidence (HC) reports in Experiment 1A (see E for Experiment 1B). For demonstration purposes, the response proportion for a given signed response offset was computed by calculating the mean response proportion across a 30° window centered around the signed response offset. (B) Violin plots of the mean response bias for HC memory reports in Experiment 1A (see F for Experiment 1B). Thick horizontal lines indicate the mean across participants. Positive values indicate memory bias towards the probe. (C) Probability distributions of signed response offsets for HC reports for ambivalent probe trials in Experiment 1A (see G for Experiment 1B). (D) Violin plots of the mean response bias for HC memory reports for ambivalent probe trials in Experiment 1A (see H for Experiment 1B).

Figure 3. Schematic of Experiment 2.
(A) Schematic of color task in Experiment 2 (B) Two perceptual comparison conditions in the color task.

Figure 4. Experiment 2 Results.
(A) Probability distribution of signed high-confidence (HC) response offsets for the shape task (see C for the color task). For demonstration purposes, the response proportion for a given signed response offset was computed by calculating the mean response proportion across a 30° window centered around the signed response offset. Positive offsets indicate memory bias toward the first similar probe. The broken lines indicate the within-subject standard errors of the mean. (B) Violin plots of the high-confidence (HC) memory bias for the shape task (see D for the color task). Thick horizontal lines indicate the mean across participants. Positive values indicate memory bias toward the first similar probe.

Figure 5. Two competing models of similarity-induced memory bias
(A) Joint density model. Noisy representation of a working memory item (X_M) is assumed to follow a von Mises distribution centered at the sample stimulus (S_M , dashed vertical line in black) with some precision (κ_M) (Left). Noisy representation of a probe item is also assumed to follow a von Mises distribution centered at the probe stimulus (S_P , dashed vertical line in blue) with some precision (κ_P) (Center). When observers decide that the memory item and the probe item are perceptually similar, then the two items are integrated to produce a joint distribution (shown in red, Right). As a result, the mean of the joint distribution is biased toward the probe item, as indicated by the dashed red vertical line. (B) Mixture density model. The initial memory representation and the probe representation are assumed to follow von Mises distributions as in the joint density model. However, this model assumes that some proportion of the memory reports are based on the probe representation. This can be accounted for by the mixture parameter

WORKING MEMORY CONTENT DISTORTED BY USE

(α). The original memory representation in this model is not biased, but the mean of the mixture distribution can be shifted toward the probe item (the dashed red vertical line).

Figure 6. The observed and simulated memory bias in Experiment 1A and 1B

(A) Probability distribution of signed human response error, best-fit joint density (JD) model responses, and best-fit mixture density model response for the shape and color tasks in Experiment 1A & 1B. (B) Correlation of bias magnitudes across the 16 probe distances between the simulated responses from the joint density model and the observed responses from human participants in the shape and color tasks in Experiment 1A and 1B. Positive values indicate a bias towards the probe item. Vertical error bars indicate ± 1 bootstrapped standard error of the simulated responses. Horizontal error bars indicate ± 1 bootstrapped standard error of human data.

Figure 7. Experiment 3 results.

Empty circles indicate individual participants. Crosses indicate participants excluded from the correlation as outliers.

Table caption

Table 1. Model comparisons.

Likelihood values for each model. Boldface indicates the preferred model. κ_p = Precision of probe representation, α = mixture parameter.