

Visual Long-Term Memory Has the Same Limit on Fidelity as Visual Working Memory

Psychological Science
24(6) 981–990
© The Author(s) 2013
Reprints and permissions:
sagepub.com/journalsPermissions.nav
DOI: 10.1177/0956797612465439
pss.sagepub.com


Timothy F. Brady¹, Talia Konkle¹, Jonathan Gill², Aude Oliva³, and George A. Alvarez¹

¹Department of Psychology, Harvard University; ²Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology; and ³Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology

Abstract

Visual long-term memory can store thousands of objects with surprising visual detail, but just how detailed are these representations, and how can one quantify this fidelity? Using the property of color as a case study, we estimated the precision of visual information in long-term memory, and compared this with the precision of the same information in working memory. Observers were shown real-world objects in random colors and were asked to recall the colors after a delay. We quantified two parameters of performance: the variability of internal representations of color (fidelity) and the probability of forgetting an object's color altogether. Surprisingly, the fidelity of color information in long-term memory was comparable to the asymptotic precision of working memory. These results suggest that long-term memory and working memory may be constrained by a common limit, such as a bound on the fidelity required to retrieve a memory representation.

Keywords

visual memory, long-term memory, short-term memory

Received 5/11/12; Revision accepted 9/26/12

A large body of work has demonstrated that visual long-term memory is capable of storing thousands of objects with significant detail (Brady, Konkle, Alvarez, & Oliva, 2008; Hollingworth, 2004, 2005; Konkle, Brady, Alvarez, & Oliva, 2010a, 2010b). However, the fidelity of long-term memory has been examined in only a qualitative way. For example, in previous work (Brady et al., 2008), we demonstrated that after seeing thousands of objects, observers succeeded at subtle object-exemplar discriminations (e.g., which of two chocolate cakes was seen) and object-state discriminations (e.g., whether the cake was half eaten or two-thirds eaten). But the information observers had to store to make these discriminations is difficult to quantify and compare across time scales and items. Therefore, several fundamental questions remain unanswered: Just how detailed is visual long-term memory? And how does its precision compare with the precision of visual working memory and of perception?

Determining the precision of long-term memory would place significant constraints on models of memory in general, and is particularly relevant for understanding the relationship between working memory and long-term memory. For example, if working memory and long-term memory have similar fidelity, then it is important to consider unified explanations for the limit on fidelity (e.g., fidelity-dependent retrieval limits), as opposed to system-specific limitations (e.g., the number of “slots” or amount of “resources” available to working memory; Wilken & Ma, 2004; Zhang & Luck, 2008). Thus, comparing working memory and long-term memory can not only help elucidate the underlying memory representation of visual

Corresponding Author:

Timothy F. Brady, Department of Psychology, Harvard University, 33 Kirkland St., Cambridge, MA 02138
E-mail: tbrady@wjh.harvard.edu

objects (Brady, Konkle, & Alvarez, 2011), but also clarify the extent to which these two stores rely on shared representations and processes (Jonides et al., 2008; McElree, 2006; Nairne, 2002).

In previous attempts to quantify the fidelity of long-term memory representations, researchers used simple stimuli, such as oriented gratings (Magnussen & Dyrnes, 1994; Magnussen, Greenlee, Aslaksen, & Kildebo, 2003). However, performance in such cases may depend on memory for decision criteria rather than on perceptual features of the objects to be remembered (Lages & Paul, 2006; Lages & Treisman, 1998; Magnussen, 2009). In addition, although oriented gratings define a well-characterized space within which to quantify fidelity, these stimuli are not suited to the strengths of long-term memory, which is best studied using meaningful stimuli, such as real-world objects (Konkle et al., 2010b). Thus, little is known about how detailed visual long-term memory representations of real-world, semantically rich objects can be.

In the study reported here, we took a psychophysical approach to quantify the fidelity of visual long-term memory for objects. We used color as a case study because the color of objects can be manipulated in a continuous space, which allowed us to extend the continuous-report paradigm used in visual working memory (Wilken & Ma, 2004) to long-term memory. Furthermore, there are metrics that allow separable, independent measurements of the fidelity of color memory and of guessing (Bays, Catalao, & Husain, 2009; Zhang & Luck, 2008). Finally, previous work has shown that continuous-report metrics do not seem to depend on verbal memory (Zhang & Luck, 2008) and that results with color generalize to shape (Zhang & Luck, 2008) and orientation (Anderson, Vogel, & Awh, 2011). Thus, we were able to quantify how accurately observers remembered a feature of a given object after seeing hundreds of objects, and how likely observers were to completely fail to retrieve a feature.

We found that from perception to working memory, observers lose significant precision in their representation of objects' color. As more items are added to working memory, the fidelity of these memory representations reaches an asymptotic limit, and, surprisingly, this limit is almost identical to the fidelity of representations in long-term memory. These results suggest that a common limit may be at work in visual working memory and long-term memory: The asymptotic fidelity observed in visual working memory may not be a consequence of a slotlike architecture (Anderson et al., 2011; Zhang & Luck, 2008, 2009) or a limited pool of resources (Bays et al., 2009; Wilken & Ma, 2004); rather, the fidelity of visual working memory may reflect a more general upper bound on how noisy a memory representation can be before it is unable to be retrieved.

Experiment 1a and 1b

In Experiment 1, observers performed a continuous-report task involving pictures of real-world objects. Observers were shown objects with randomly selected hues and asked to choose from a color wheel what hues the objects were. Such continuous-report methods have been used to investigate working memory for simple geometric shapes (e.g., Brady & Alvarez, 2011; Wilken & Ma, 2004; Zhang & Luck, 2008), but have never been adapted for examining long-term memory.

This method allowed us to measure the fidelity of perception, working memory, and long-term memory using a within-subjects design. In the *perception* condition, observers had to match the color of a visible object. In the *working memory* condition, observers were given 3 s to encode three objects and then had to report the color of each object after a 1-s delay. We used three objects per trial to match the set size at which working memory fidelity reaches asymptote (Anderson et al., 2011; Zhang & Luck, 2008). In the *long-term memory* condition, observers viewed hundreds of objects, presented one at a time, and then were asked to report the color of every single object, one at a time. In Experiment 1a, observers had 3 s to encode each object in the long-term memory condition, so that the total time that a display was visible in this condition matched the total time that a display was visible in the working memory condition. In Experiment 1b, we gave observers only 1 s to encode each object in the long-term memory condition, so that the display time per object in this condition matched the display time per object in the working memory condition. Observers saw different objects in each of the three conditions.

Method

Participants. Fourteen observers (age range = 18–25 years) participated in Experiment 1: 5 in Experiment 1a and 9 in Experiment 1b. They gave informed consent and had normal color vision (assessed using Ishihara's, 1936, test for color deficiencies). All participants completed all three conditions, with the order randomized across participants.

Apparatus. Experiments were run in MATLAB using the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997). Stimuli were presented on a 49° × 31° display and viewed from a distance of 57 cm.

Stimuli. Five hundred forty pictures of categorically distinct objects were selected from Brady et al. (2008).¹ We chose objects that were largely in a single arbitrary color (e.g., each object would be recognizable in any color; see Fig. 1). Objects were rotated randomly in hue space, such

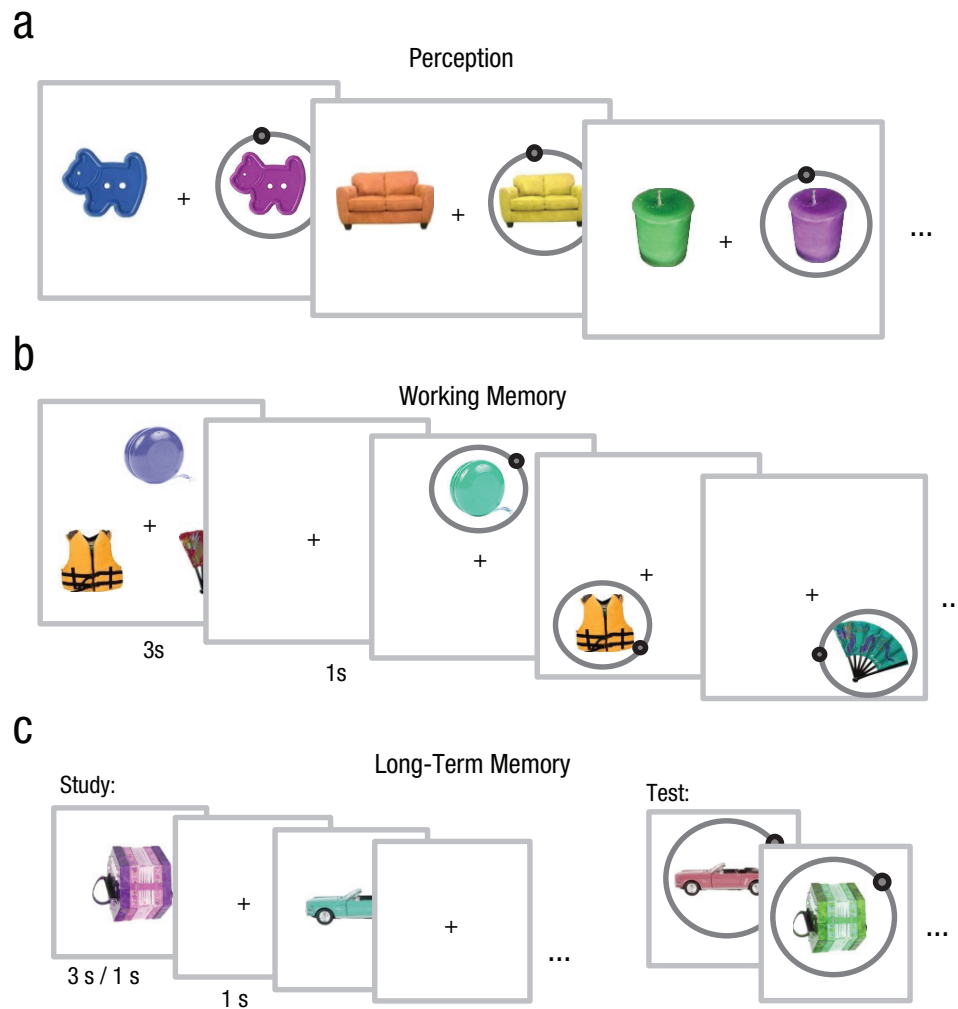


Fig. 1. Illustration of the method of Experiment 1. Observers were shown objects that had colors randomly rotated in hue and were asked to report each object's color (a) while it was still visible (perception condition), (b) after a 1-s delay (working memory condition), or (c) after seeing several hundred objects over 30 min (long-term memory condition). In Experiment 1a, each item in the long-term memory condition was shown for 3 s; in Experiment 1b, each item in that condition was shown for 1 s.

that on each trial, the initial object color was determined by adding a random angle between 0° and 359° to the original hue. Across observers, each image appeared equally often in the three conditions. All stimuli subtended approximately 6° of visual angle.

Perception condition. To assess the fidelity of color perception, we had participants perform a color-matching task (see Fig. 1a). On each of the 180 trials, two copies of the same image were presented simultaneously, centered 5° to the left and right of fixation. The left image was the *standard* image, and the right one was the *test* image (initially presented in gray scale). The task was to adjust the color of the test item to match the standard.

Working memory condition. On each trial, three objects were presented simultaneously for 3 s, in a circle around fixation (see Fig. 1b). Participants were instructed to remember the color of all three objects. The objects disappeared for 1 s, and then memory for the color of the items was tested one at a time in a randomly chosen sequence. Participants completed 60 trials, for a total of 180 tests.

Long-term memory condition. During the study block, participants viewed 232 images presented one at a time, for either 1 s (Experiment 1a) or 3 s (Experiment 1b) each (see Fig. 1c). There was a 1-s blank between images. Participants were instructed to remember the color and

identity of each object as they viewed the images. During this block, participants performed a repeat-detection task intended to encourage them to maintain focused attention. Twenty-six images in the study stream appeared twice in a row, and participants pushed the space bar when they noticed a repeat. Participants were given feedback only when they responded; a red fixation cross indicated an incorrect response, and a green fixation cross indicated a correctly detected repeat. No feedback was given for misses or correct rejections.

Immediately after the study block, we tested the fidelity with which participants remembered the color of the objects. Items that were repeated in the study stream were not tested, so there were 180 tested images.

Continuous report. In each of the three conditions, participants' color memory was measured using the method of adjustment. At the beginning of each test, the item appeared in gray scale, with the mouse pointer at the center of the item. When the participant moved the mouse, the test item appeared in color. The angle between the mouse and the center of the test item determined the item's hue, and a dot presented along an adjustment ring surrounding the item indicated the current angular position of the mouse. When participants decided that the current color was correct, they clicked the mouse. The angular error was taken as a measure of accuracy. The color wheel was randomly rotated across trials.

Participants proceeded at their own pace and were asked to be as accurate as possible in their decisions. Feedback was given after accurate responses: The words "good," "great," or "perfect" appeared on the screen for errors of less than 10°, 5°, or 0°, respectively.

Data analysis. On any given trial, we measured error in degrees, between 0° (perfect memory) and ±180° (poor memory). In the continuous-report paradigm, the histogram of errors over trials typically shows that responses are centered around 0°, but that across all responses, there are errors distributed across the entire range. The error histograms we obtained for the three conditions (see Fig. 2) were well fit by a mixture of two distributions: (a) a Gaussian-like distribution (defined on a circular space as a von Mises distribution), taken to reflect successful memory retrieval with some degree of precision, and (b) a uniform distribution, taken to reflect random guessing (Zhang & Luck, 2008). We used Zhang and Luck's method to separate trials in which the color was retrieved with some level of fidelity and trials in which the color of the item was forgotten.

The fidelity (precision) of memory representation was estimated as the standard deviation of the von Mises

distribution. The narrower the distribution around 0°, the more precise the memory representation. The probability of guessing was estimated by the height of the uniform distribution. Maximum likelihood estimation was used to estimate these two parameters for each condition.

Results

Figure 2 shows the distribution of errors in the perception, working memory, and long-term memory conditions, combined across Experiment 1a and Experiment 1b. The fidelity and guessing parameters for each experiment are summarized in Figure 3.

Experiment 1a. In the perception condition, observers were highly accurate, with precision estimated at 6.7° ($SEM = 0.8$) and the probability of guessing estimated at .0 ($SEM = .0$). Thus, when the stimulus was present on the screen, observers never responded randomly and had a tight distribution centered on the correct color.

Results in the working memory condition were in line with Zhang and Luck's (2008) findings. Observers' precision was 19.0° ($SEM = 1.3$), and their probability of guessing was .09 ($SEM = .02$). Thus, there was a major change in fidelity from perception to working memory: The standard deviation increased by 183%, a serious cost in memory fidelity for having to hold the items in mind for several seconds, $t(4) = 10.5$, $p < .0001$.

In the long-term memory condition (3 s/item), observers' precision was 20.3° ($SEM = 3.3$), and the probability of guessing was .58 ($SEM = .05$). The increase in the guess rate from the working memory condition was quite large (from .09 to .58), $t(4) = 17.5$, $p < .0001$. However, surprisingly, the fidelity observed for 180 items in long-term memory was not significantly different from the fidelity observed for 3 items at a time in working memory, $t(4) = -0.81$, $p = .46$. Note that the estimated precision of working memory in this experiment is similar to the precision observed in several working memory experiments that tested memory for color patches at set sizes of 3 and greater (Zhang & Luck, 2008, 2011), even though we used real-world objects.

Experiment 1b. In this experiment, observers had only 1 s to encode the color of each item in the long-term memory condition. Despite this severe decrease in encoding time, Experiment 1b replicated Experiment 1a nearly exactly (see Fig. 3). Fidelity was 4.7° ($SEM = 0.5$) in the perception condition, 17.8° ($SEM = 1.0$) in the working memory condition, and 19.3° ($SEM = 0.9$) in the long-term memory condition. The probabilities of guessing were .006 ($SEM = .002$), .08 ($SEM = .01$), and .63 ($SEM = .05$), respectively. As before, the fidelity of working

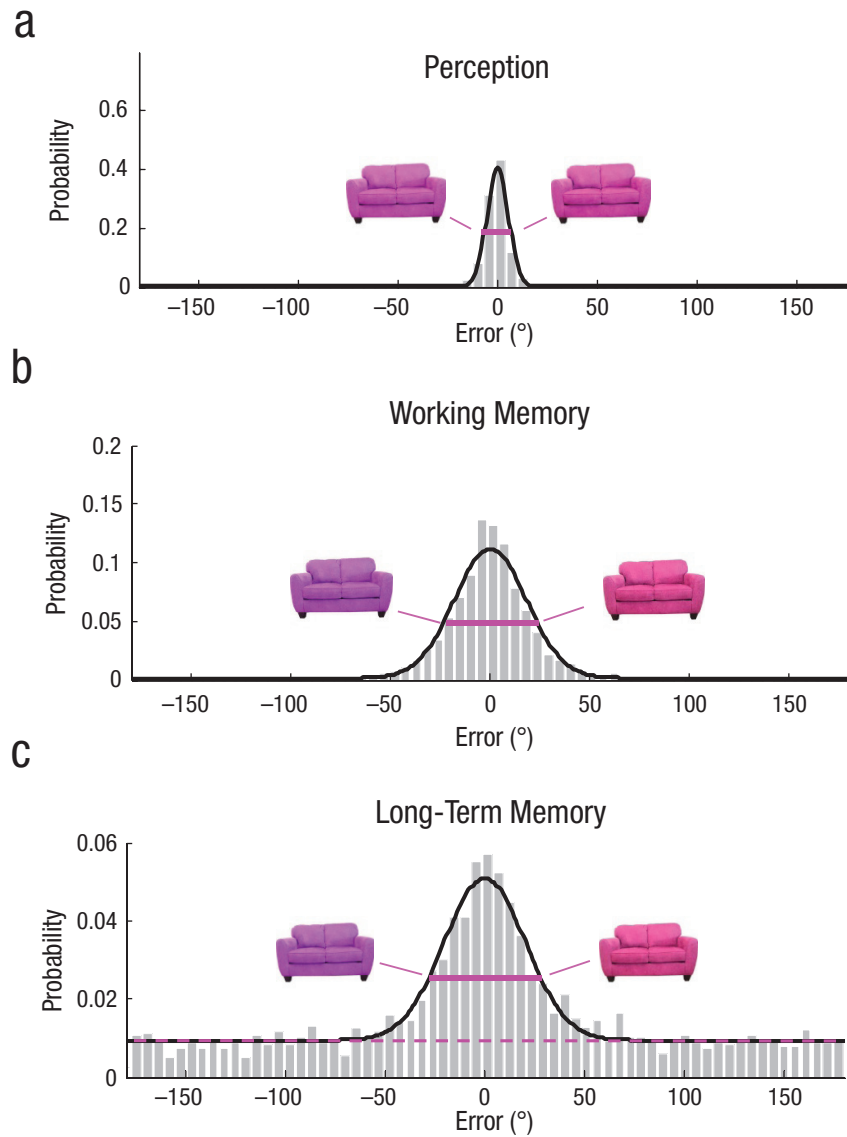


Fig. 2. Results pooled across all observers in Experiments 1a and 1b. The histograms represent the distribution of the magnitude of error in observers' responses in the (a) perception, (b) working memory, and (c) long-term memory conditions. The black curves show the model fits from a mixture model that combines a uniform guessing distribution with a Gaussian-like distribution of correct responses. The pink solid lines show the width of the Gaussian at 1 standard deviation and are flanked by illustrations showing the corresponding colors (± 1 SD of error) from a sample trial with a picture of a couch. The pink dashed lines show the guessing distributions alone, without the Gaussian component.

memory and the fidelity of long-term memory were not significantly different, $t(8) = 1.0$, $p = .33$. In addition, there was no significant difference in the fidelity of long-term memory between the two experiments, $t(12) = 0.36$, $p = .72$. Thus, the extra encoding time in Experiment 1a made no difference to the fidelity of color information in long-term memory. When we combined results across Experiments 1a and 1b, we again did not find a significant difference between the fidelity of working memory

and the fidelity of long-term memory ($M_s = 18.2^\circ$ and 19.7°), $t(13) = 0.84$, $p = .41$.

Discussion

We measured the fidelity of color information in visual long-term memory in two studies and compared it with the fidelity of working memory and perception. The data show an extremely precise fidelity in perception ($\sim 5\text{--}6^\circ$),

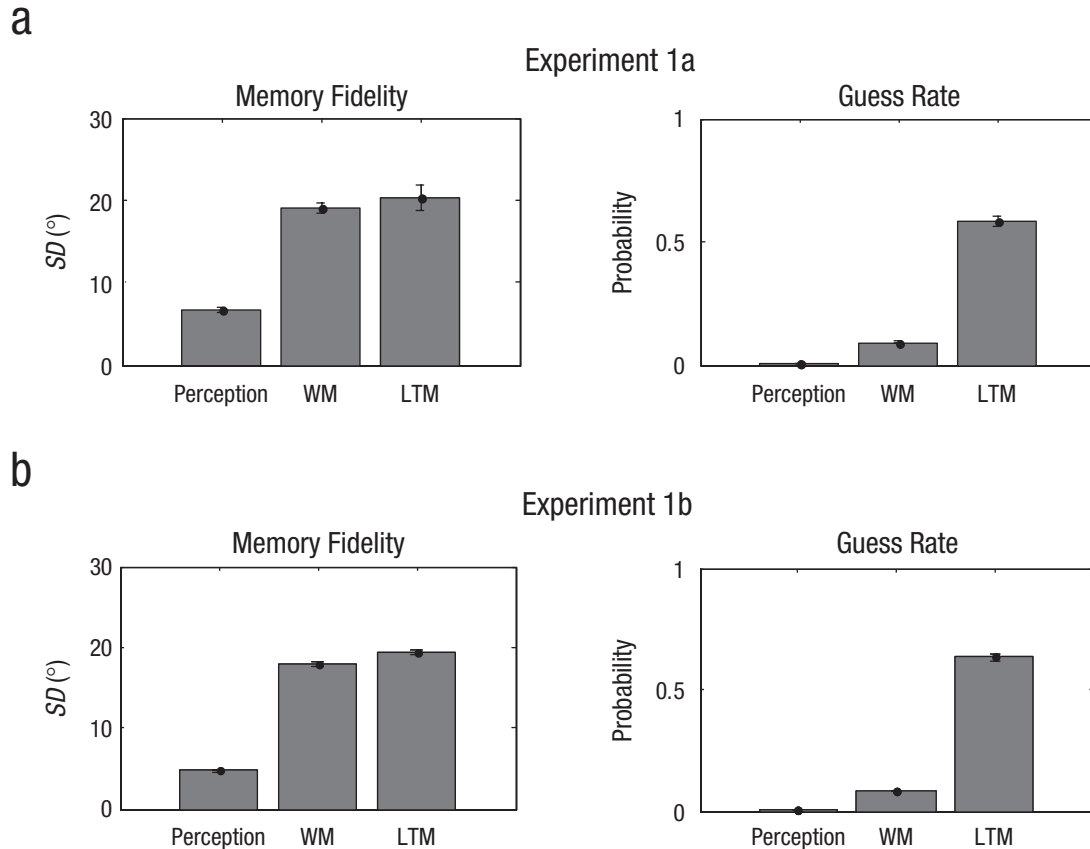


Fig. 3. Estimated fidelity (standard deviation of the von Mises distribution) and probability of guessing in (a) Experiment 1a and (b) Experiment 1b. Results are shown for each of the three conditions: perception, working memory (WM), and long-term memory (LTM). Error bars represent standard errors of the mean.

but a significantly lower fidelity in both working memory and long-term memory. Surprisingly, the fidelity for color was comparable between working memory and long-term memory ($\sim 20^\circ$). This was true both when long-term memory encoding time matched the total encoding time in the working memory condition (Experiment 1a) and when long-term memory encoding time matched the per-item encoding time in the working memory condition (Experiment 1b). The results showed that nearly all of the information loss from working memory to long-term memory is accounted for by an increased chance of entirely losing an item's color from memory (increased guess rate).

In the long-term memory condition, participants had to store hundreds of items for long durations and were required to encode and then retrieve the items, whereas in the working memory condition, participants could keep items and their colors actively in mind. Despite these major differences between the two tasks, the fidelity of working memory (when three items were held in mind) and the fidelity of long-term memory were nearly

identical. This indicates that observers have highly detailed long-term memory representations—even when fidelity is measured quantitatively rather than with qualitative forced-choice comparisons (Brady et al., 2008; Hollingworth, 2004).

Experiment 2

It is possible that long-term memory and working memory have the same fidelity because long-term memory representations inherit their fidelity directly from working memory. For example, if items have to enter working memory to be encoded into long-term memory, and if there is no further degradation of representations once they are encoded, then long-term memory representations would have exactly the same fidelity as working memory representations. Although this is a possible account of our results, previous work has shown that the fidelity of visual working memory depends on the number of items remembered (Wilken & Ma, 2004; Zhang & Luck, 2008). Thus, because items in our long-term

memory task were presented sequentially, one at a time, this *inherited-precision* account predicts that the precision of memory representations in that task should match the fidelity of working memory for a single item. Results from Experiment 1 cannot directly address this prediction, because multiple items were presented simultaneously in the working memory task.

To test this inherited-precision hypothesis, we matched the encoding conditions in the working memory and long-term memory tasks in Experiment 2. To preview the results, we found that observers can in fact remember a single item with better precision in working memory than in long-term memory, which is inconsistent with the inherited-precision account.

Method

Participants. Six observers participated in Experiment 2. None had participated in Experiment 1. All participants gave informed consent, were between the ages of 18 and 25, and had normal color vision (assessed using Ishihara's test for color deficiencies).

Stimuli. The stimuli were the same objects as in Experiment 1.

Procedure. The procedure was identical to that of Experiment 1a, except that the working memory condition was modified to consist of 180 trials with only a single item presented on each trial. Each object was presented for 3 s and then tested after a 1-s delay.

Results

Results for the perception condition (fidelity = 5.6° , $SEM = 1.1$; probability of guessing = .01, $SEM = .01$) and the long-term memory condition (fidelity = 20.5° , $SEM = 7.0$; probability of guessing = .67, $SEM = .15$) replicated the results from Experiments 1a and 1b. However, working memory fidelity for one item (14.5° , $SEM = 1.3$) was significantly better than long-term memory fidelity, given matched encoding conditions, $t(5) = 2.96$, $p = .03$. In addition, comparing results across experiments revealed that the fidelity of working memory for one item was significantly better than the fidelity of working memory for three items—comparison with Experiment 1a: $t(9) = 2.45$, $p = .03$; comparison with Experiment 1b: $t(13) = 2.06$, $p = .06$. These results show that the memory precision of color information of real-world object stimuli is not fixed at encoding; it is possible for the fidelity of working memory to be better than the fidelity of long-term memory.

Discussion

The results of Experiment 2 show that when a single real-world object is encoded, the fidelity of its representation in

working memory is higher than the fidelity the representation will have when it is later probed in long-term memory. This indicates that the fidelity of long-term memory is not directly inherited from working memory, and that there is additional degradation in long-term memory that reduces the measured precision of retrieved items to 20° .

Intriguingly, the data across all experiments show that the level at which the fidelity of working memory plateaus is identical to the fidelity observed in long-term memory (Fig. 4). To bolster this finding, we conducted several control experiments (see also the Supplemental Material available online).

First, we tested working memory using displays with five items (Control Experiment 1), to more clearly demonstrate the plateau in the fidelity of working memory (see the green line in Fig. 4). Next, we asked whether we could make long-term memory precision worse than this limit. We reasoned that if we doubled the number of items in memory from 180 to 360, this might lead to less precise memories. However, we instead found that this manipulation only increased the probability of guessing, and the fidelity of the remembered items remained at about 20° (Control Experiment 2). Finally, we examined whether long-term memory precision could be more precise than this limit—which would still be consistent with a limit on the fidelity of the memory representations. However, surprisingly, we found that even when the memory set consisted of only 20 items (Control Experiment 3), precision was similar; although there was a benefit in overall performance relative to when the memory set included 180 or 360 items, this benefit was due to a lowered guessing rate. The same level of precision was found when participants performed a verbal interference task (Control Experiment 4), a result suggesting that the fidelity limit is not likely due to a verbal coding strategy.

In summary, across a very wide range of overall difficulty levels in the long-term memory task, with guess rates ranging from .26 to .73, we found that fidelity remained constant at a standard deviation of 19° to 20° . Any representation more variable than this limit seems to be lost entirely to guessing, in both working memory and long-term memory.

General Discussion

Across several experiments, we found that observers lose significant precision in their representation of real-world objects when going from perception to working memory. However, the precision of three or four actively maintained representations in working memory is the same as that of hundreds of representations encoded and then retrieved from long-term memory. Thus, long-term memory fidelity is significantly higher than previously believed, even when quantified using psychophysical methods.

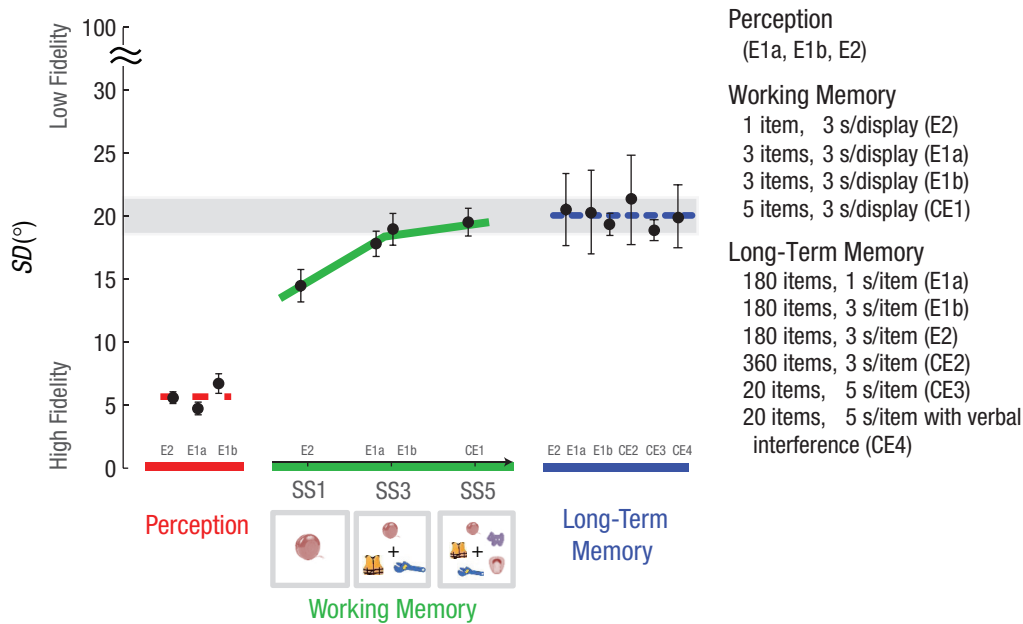


Fig. 4. Fidelity (standard deviation of the von Mises distribution) estimates from all experiments. Perceptual fidelity is plotted for Experiments 1a, 1b, and 2. Working memory fidelity is plotted for Experiments 1a and 1b (three-item displays), Experiment 2 (one-item displays), and Control Experiment 1 (five-item displays). Long-term memory fidelity is plotted for Experiments 1a, 1b, and 2; a control experiment with 360 items rather than 180 (Control Experiment 2); a control experiment with only 20 items and a longer encoding time (Control Experiment 3); and a control experiment with verbal interference (Control Experiment 4). The solid line for working memory shows fidelity as a function of set size. The dashed lines for perception and long-term memory indicate the mean standard deviation for those conditions. Error bars represent standard errors of the mean. The list to the right of the graph summarizes key characteristics of the various experiments (see the Supplemental Material available online for additional details about the control experiments). The gray rectangle highlights the fidelity at which working memory reaches a plateau, which is the same as the fidelity estimated for visual long-term memory. CE = control experiment; E = experiment; SS = set size.

Furthermore, the fidelity of long-term memory is not directly inherited from the fidelity of representations at encoding, but instead seems to represent an asymptotic limit on the fidelity of items retrieved from memory: In working memory, as the number of items stored increases, fidelity plateaus at a standard deviation of about 20° . Similarly, fidelity of items retrieved from long-term memory is degraded relative to items retrieved from working memory, but for remembered items, fidelity of long-term memory does not get worse than a standard deviation of approximately 20° , despite the necessity of representing more items for a longer duration and making use of an encoding and retrieval process rather than active storage. Additionally, this degree of fidelity for long-term memory is robust to a variety of encoding durations, a variety of number of objects to be stored, and the presence or absence of a verbal interference task. Thus, we suggest that a standard deviation of about 20° may represent a limit on the fidelity of arbitrary color information that can be successfully retrieved from memory: Any memory representations that degrade so that they have more variability than a standard deviation of about 20° seem to be

irretrievable (for a visualization of this memory fidelity, see Fig. 5).

This pattern of results suggests a dramatic reinterpretation of existing data on working memory: The plateau in working memory fidelity is likely not caused by factors intrinsic to working memory, such as the fidelity of a slot (Anderson et al., 2011; Zhang & Luck, 2008, 2009) or the quantity of a resource (Bays et al., 2009), but is instead a general property of the memory encoding and retrieval system. That is, the fidelity of working memory and long-term memory may reflect an upper bound on how noisy a memory representation can be before it is unable to be retrieved.

Relationship between working memory and long-term memory

Several influential studies have found that working memory fidelity plateaus at a standard deviation of approximately 20° (Anderson & Awh, 2012; Anderson et al., 2011; Zhang & Luck, 2008, 2009, 2011). In particular, fidelity does not seem to decrease when more than three

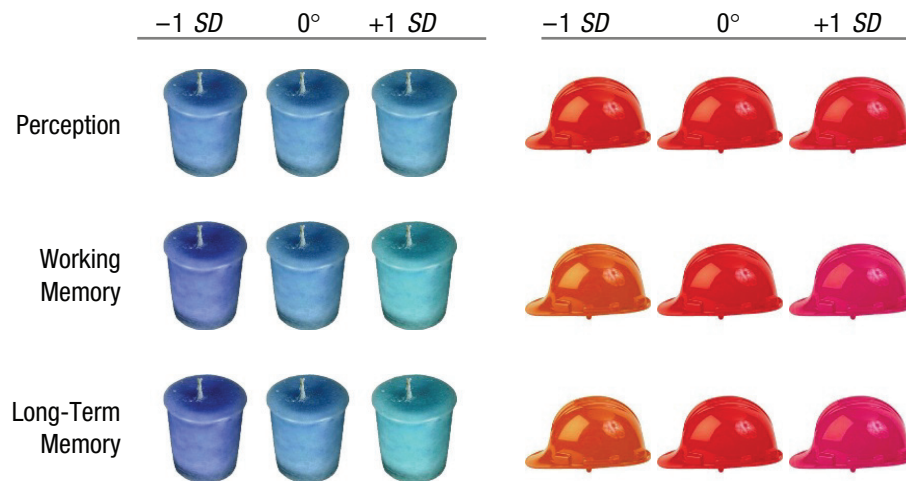


Fig. 5. Pictorial representation of the memory fidelity observed in perception, working memory, and long-term memory. In each triplet, the central stimulus shows the studied color, and the items to the right and left represent colors 1 standard deviation above and below the studied color ($\pm 6^\circ$ in perception, $\pm 19^\circ$ in working memory for three items, $\pm 20^\circ$ in long-term memory).

or four items are encoded (Zhang & Luck, 2008) or when observers hold items for longer durations (Zhang & Luck, 2009). On the basis of this apparent asymptote in fidelity, all of these researchers have concluded that working memory represents items with discrete slots that undergo catastrophic failures when items are held for long durations (Zhang & Luck, 2009).

However, these explanations for why fidelity does not become worse are based entirely on models of active storage in working memory (slots, resources). For example, Zhang and Luck (2008) interpreted this asymptote as resulting from a limited number of memory slots that maintain fixed-resolution representations in working memory. According to this theory, if you have three slots in memory, you can use them to represent fewer than three items with more precision by allocating multiple slots per item, but after you have three items in memory, you can no longer split your representations among more items, so all subsequent items are simply not encoded. This failure to encode more than three items results in an increased probability of guessing but a flat fidelity asymptote as the number of items exceeds the number of slots. Other researchers have argued that the asymptote is a natural consequence of spreading a continuously divisible memory resource across multiple items, which leads to decreased precision and an increased likelihood of forgetting items as set size increases (e.g., Bays & Husain, 2008; Wilken & Ma, 2004).

However, our finding that long-term memory shares a similar limit suggests an alternative model. Rather than reflecting an intrinsic property of the working memory system, this asymptotic fidelity limit may instead reflect a

property of the broader memory system and factors that limit memory retrieval.

Conclusion: shared limits for an integrated visual memory system

The broader working memory literature—particularly the literature on verbal stimuli—has accumulated significant evidence for shared principles between short-term and long-term memory (Jonides et al., 2008; McElree, 2006; Nairne, 2002). For example, items putatively held in active storage are not accessed any faster than those held in passive storage (McElree, 2006). In addition, a number of empirical results highlight that working memory tasks do not isolate working memory mechanisms independently of long-term mechanisms. For example, performance on any given working memory trial is influenced by previous trials, an influence that includes systematically induced biases and proactive interference (Hartshorne, 2008; Huang & Sekuler, 2010; Makovski & Jiang, 2008). These findings suggest an obligatory influence of long-term storage on working memory (Brady et al., 2011; see also Olson, Moore, Stark, & Chatterjee, 2006, for evidence from neuroscience).

The present empirical results showing that long-term memory fidelity is so high—and, in fact, equivalent to the asymptotic fidelity of working memory—lead us to propose a new link between working memory and long-term memory: They appear to have the same lower bound on memory fidelity. Recalled items are never noisier than a fixed limit; after that limit, an item is lost, perhaps irretrievable via conscious access because it no

longer sufficiently resembles the original memory trace that was laid down.

Declaration of Conflicting Interests

The authors declared that they had no conflicts of interest with respect to their authorship or the publication of this article.

Funding

This work was supported in part by the National Science Foundation under Grant No. 1016862 to A. O. and by a faculty research award from Google to A. O.

Supplemental Material

Additional supporting information may be found at <http://pss.sagepub.com/content/by/supplemental-data>

Note

1. The stimuli may be found on Timothy F. Brady's Web site, <http://timbrady.org>

References

- Anderson, D. E., & Awh, E. (2012). The plateau in mnemonic resolution across large set sizes indicates discrete resource limits in visual working memory. *Attention, Perception, & Psychophysics*, *74*, 891–910.
- Anderson, D. E., Vogel, E. K., & Awh, E. (2011). Precision in visual working memory reaches a stable plateau when individual item limits are exceeded. *Journal of Neuroscience*, *31*, 1128–1138.
- Bays, P. M., Catalao, R. F. G., & Husain, M. (2009). The precision of visual working memory is set by allocation of a shared resource. *Journal of Vision*, *9*(10), Article 7. Retrieved from <http://www.journalofvision.org/content/9/10/7>
- Bays, P. M., & Husain, M. (2008). Dynamic shifts of limited working memory resources in human vision. *Science*, *321*, 851–854.
- Brady, T. F., & Alvarez, G. A. (2011). Hierarchical encoding in visual working memory: Ensemble statistics bias memory for individual items. *Psychological Science*, *22*, 384–392.
- Brady, T. F., Konkle, T., & Alvarez, G. A. (2011). A review of visual memory capacity: Beyond individual items and toward structured representations. *Journal of Vision*, *11*(5), Article 4. Retrieved from <http://www.journalofvision.org/content/11/5/4>
- Brady, T. F., Konkle, T., Alvarez, G. A., & Oliva, A. (2008). Visual long-term memory has a massive storage capacity for object details. *Proceedings of the National Academy of Sciences, USA*, *105*, 14325–14329.
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*, 433–436.
- Hartshorne, J. K. (2008). Visual working memory capacity and proactive interference. *PLoS ONE*, *3*(7), e2716. Retrieved from <http://www.plosone.org/article/lookup?articleURI=info:doi/10.1371/journal.pone.0002716>
- Hollingworth, A. (2004). Constructing visual representations of natural scenes: The roles of short- and long-term visual memory. *Journal of Experimental Psychology: Human Perception and Performance*, *30*, 519–537.
- Hollingworth, A. (2005). The relationship between online visual representation of a scene and long-term scene memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *31*, 396–411.
- Huang, J., & Sekuler, R. (2010). Distortions in recall from visual memory: Two classes of attractors at work. *Journal of Vision*, *10*(2), Article 24. Retrieved from <http://www.journalofvision.org/content/10/2/24>
- Ishihara, S. (1936). *The series of plates designed as tests for colour-blindness*. Tokyo, Japan: Kanehara & Co.
- Jonides, J., Lewis, R. L., Nee, D. E., Lustig, C. A., Berman, M. G., & Moore, K. S. (2008). The mind and brain of short-term memory. *Annual Review of Psychology*, *59*, 193–224.
- Konkle, T., Brady, T. F., Alvarez, G. A., & Oliva, A. (2010a). Conceptual distinctiveness supports detailed visual long-term memory for real-world objects. *Journal of Experimental Psychology: General*, *139*, 558–578.
- Konkle, T., Brady, T. F., Alvarez, G. A., & Oliva, A. (2010b). Scene memory is more detailed than you think: The role of categories in visual long-term memory. *Psychological Science*, *21*, 1551–1556.
- Lages, M., & Paul, A. (2006). Visual long-term memory for spatial frequency? *Psychonomic Bulletin & Review*, *13*, 486–492.
- Lages, M., & Treisman, M. (1998). Spatial frequency discrimination: Visual long-term memory or criterion setting? *Vision Research*, *38*, 557–572.
- Magnussen, S. (2009). Implicit visual working memory. *Scandinavian Journal of Psychology*, *50*, 535–542.
- Magnussen, S., & Dyrnes, S. (1994). High-fidelity perceptual long-term memory. *Psychological Science*, *5*, 99–102.
- Magnussen, S., Greenlee, M. W., Aslaksen, P. M., & Kildebo, O. Ø. (2003). High-fidelity perceptual long-term memory revisited—and confirmed. *Psychological Science*, *14*, 74–76.
- Makovski, T., & Jiang, Y. V. (2008). Proactive interference from items previously stored in visual working memory. *Memory & Cognition*, *36*, 43–52.
- McElree, B. (2006). Accessing recent events. In B. H. Ross (Ed.), *The psychology of learning and motivation* (Vol. 46, pp. 155–200). San Diego, CA: Academic Press.
- Nairne, J. S. (2002). Remembering over the short-term: The case against the standard model. *Annual Review of Psychology*, *53*, 53–81.
- Olson, I. R., Moore, K. S., Stark, M., & Chatterjee, A. (2006). Visual working memory is impaired when the medial temporal lobe is damaged. *Journal of Cognitive Neuroscience*, *18*, 1087–1097.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*, 437–442.
- Wilken, P., & Ma, W. J. (2004). A detection theory account of change detection. *Journal of Vision*, *4*(12), Article 11. Retrieved from <http://www.journalofvision.org/content/4/12/11>
- Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representations in visual working memory. *Nature*, *452*, 233–235.
- Zhang, W., & Luck, S. J. (2009). Sudden death and gradual decay in visual working memory. *Psychological Science*, *20*, 423–428.
- Zhang, W., & Luck, S. J. (2011). The number and quality of representations in working memory. *Psychological Science*, *22*, 1434–1441.

SUPPLEMENTAL ONLINE MATERIAL

Control Experiment 1: 5 items in working memory.

This experiment was designed to verify the plateau in working memory fidelity was present with our real-world object displays. This experiment was identical to Experiment 1a's working memory condition except five items were presented in working memory rather than three. All timing and feedback was the same as Experiment 1a. The average probability of guessing was 0.28 ($SEM = 0.01$), and observers' precision was 19.5° ($SEM = 1.1$).

Control Experiment 2: 360 items in long-term memory.

This experiment was designed to test whether long-term memory might be less precise when more items were studied. This experiment was identical to Experiment 1a's long-term memory condition, except that observers studied 360 objects rather than 180. All timing and feedback was the same as Experiment 1a. The average probability of guessing was 0.73 ($SEM = 0.03$), and observers' precision was 19.6° ($SEM = 2.3$).

Control Experiment 3: 20 items at a time in long-term memory.

This experiment was designed to test whether long-term memory might be more precise when fewer items were studied; when these items were presented for longer durations; and when the length of time they needed to be held in memory for was shorter. Items were presented one at a time for 5s each and observers were tested after blocks of 20 objects. In between studying the 20 objects and being tested on them, observers had to perform two trials of a color change detection task to ensure we were tapping long-term memory rather than working memory. This change detection task was identical to that of

Brady and Tenenbaum (2013)'s Experiment 2. Observers performed 10 repetitions of 20-object study blocks, each consisting of new objects, for a total of 200 memory trials. As expected, overall performance was quite high: the average probability of guessing was only 0.29 ($SEM = 0.05$). However, despite this high overall performance, the precision of observers' reports was nearly identical to Experiments 1a, 1b and 2 (18.9° ; $SEM = 0.83$).

Control Experiment 4: Verbal interference in long-term memory.

This experiment was designed to examine the role verbal labeling might play in the long-term memory experiments. This experiment was identical to Control Experiment 3, but all items were studied under conditions of verbal interference. Observers rehearsed 4 randomly-generated digits aloud and we monitored them to ensure that they were continuously speaking these digits throughout the encoding period (a common verbal interference task for working memory studies; e.g., Hollingworth, 2005). These digits were randomly generated and displayed to observers at the beginning of each block. This task ensures that observers cannot be using verbal memory to encode the colors (e.g., "it's a pink couch") which could artificially lead both working memory and long-term memory to appear to have a constant fidelity limit. With this dual task paradigm, we found a slightly higher probability of guessing (0.42, $SEM = 0.11$) but no change in precision (20° , $SEM = 2.5$). This provides evidence that the fidelity of representations in visual long-term memory is not strongly influenced by the ability to verbally encode the color names.

References

- Brady, T. F., and Tenenbaum, J.B. (2013). A probabilistic model of visual working memory: Incorporating higher-order regularities into working memory capacity estimates. *Psychological Review*, 120(1), 85-109.
- Hollingworth, A. (2005). The relationship between online visual representation of a scene and long-term scene memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31, 396-411.