Journal of Experimental Psychology: Human Perception and Performance

Evidence of Gradual Loss of Precision for Simple Features and Complex Objects in Visual Working Memory

Rosanne L. Rademaker, Young Eun Park, Alexander T. Sack, and Frank Tong

Online First Publication, March 1, 2018. http://dx.doi.org/10.1037/xhp0000491

Journal of Experimental Psychology:
Human Perception and Performance

0096-1523/18/$12.00 http://dx.doi.org/10.1037/xhp0000491

Evidence of Gradual Loss of Precision for Simple Features and Complex Objects in Visual Working Memory

Rosanne L. Rademaker
University of California San Diego and Maastricht University

Young Eun Park
Vanderbilt University

Alexander T. Sack
Maastricht University

Frank Tong
Vanderbilt University

Previous studies have suggested that people can maintain prioritized items in visual working memory for many seconds, with negligible loss of information over time. Such findings imply that working memory representations are robust to the potential contaminating effects of internal noise. However, once visual information is encoded into working memory, one might expect it to inevitably begin degrading over time, as this actively maintained information is no longer tethered to the original perceptual input. Here, we examined this issue by evaluating working memory for single central presentations of an oriented grating, color patch, or face stimulus, across a range of delay periods (1, 3, 6, or 12 s). We applied a mixture-model analysis to distinguish changes in memory precision over time from changes in the frequency of outlier responses that resemble random guesses. For all 3 types of stimuli, participants exhibited a clear and consistent decline in the precision of working memory as a function of temporal delay, as well as a modest increase in guessing-related responses for colored patches and face stimuli. We observed a similar loss of precision over time while controlling for temporal distinctiveness. Our results demonstrate that visual working memory is far from lossless: while basic visual features and complex objects can be maintained in a quite stable manner over time, these representations are still subject to noise accumulation and complete termination.

Public Significance Statement

The ability to retain visual information over brief delays is critical for accurate visual performance. Numerous studies have claimed that items in visual working memory are immune to temporal decay. This is surprising, given that information maintained in working memory is no longer tethered to the original perceptual input. One could rightfully wonder how any biological system would be able to achieve perfect retention under such circumstances. Here, we show that working memory for individual orientations, colors, and faces, undergoes gradual decay over time, even when occurrences of complete memory failure are taken into account. Our statistical and model comparison analyses provide compelling evidence that information represented in visual working memory accumulates noise over time, leading to a gradual but inevitable loss of visual precision.

Keywords: visual working memory, orientation perception, color perception, face perception, memory decay

Supplemental materials: http://dx.doi.org/10.1037/xhp0000491.supp

Rosanne L. Rademaker, Department of Psychology, University of California San Diego, and Cognitive Neuroscience Department, Maastricht University; Young Eun Park, Psychology Department, Vanderbilt University; Alexander T. Sack, Cognitive Neuroscience Department, Maastricht University; Frank Tong, Psychology Department, Vanderbilt University.

Rosanne L. Rademaker and Young Eun Park contributed equally to this article.

This research was supported by National Science Foundation Grant BCS-1228526 to Frank Tong and facilitated by a National Institutes of Health P30-EY-008126 Center Grant to the Vanderbilt Vision Research Center. We thank J. Serences for graciously providing the lab space and resources that were necessary to conduct Experiment 1B.

Data and Stimulus code are available on the Open Science Forum (OSF) at https://osf.io/jmkc9/.

Correspondence concerning this article should be addressed to Frank Tong, Psychology Department, Vanderbilt University, 111 21st Ave South, 301 Wilson Hall, Nashville, TN 37240 or to Rosanne L. Rademaker, Department of Psychology, UC San Diego, 9500 Gilman Dr, La Jolla, CA 92093. E-mail: frank.tong@vanderbilt.edu or rrademaker@ucsd.edu

Information about a recently viewed visual object can be actively maintained to fulfill perceptual and cognitive goals, such as when comparing one visual item to another, or when planning a series of visually guided actions (Hayhoe, Aivar, Shrivastavah, & Mruczek, 2002; Hayhoe, Shrivastava, Mruczek, & Pelz, 2003; Hollingworth, Richard, & Luck, 2008). To perform these tasks, visual information must be stored with high fidelity and protected against interference and the effects of time. Visual working memory serves to support these functions, allowing for the flexible and active maintenance of precise visual information over brief delays to provide online support for one's immediate cognitive goals.

A seemingly simple issue, regarding the fate of visual memories as they evolve over time has proved equivocal. Intuitively, one might expect that once a stimulus is gone from view, that stimulus cannot be maintained perfectly by the mind alone—without direct perception to keep it grounded. Nevertheless, studies assessing visual memory over time have generally demonstrated its high fidelity and robustness to decay over time.

In earlier work, researchers applied psychophysical procedures to measure discrimination thresholds for simple visual features, varying the temporal delay between the two items to be compared. For example, Regan and Beverley (1985) found negligible loss of visual precision in the ability to discriminate between pairs of orientations presented either 1 s apart or 10 s apart, implying that visually precise information can be sustained by the central nervous system in a lossless manner over this interval. Many subsequent studies have made similar reports of finding no loss of visual resolution over time when assessing working memory for spatial frequency (Bennett & Cortese, 1996; Huang & Sekuler, 2010; Magnussen, Greenlee, Asplund, & Dyrnes, 1991; Regan, 1985), motion direction (Blake, Cepeda, & Hiris, 1997), motion velocity (Magnussen & Greenlee, 1992), and even memory for complex stimuli such as faces (Banko, Gal, & Vidnyanszky, 2009). Although these claims are based on reporting a null difference, they do support the notion that visual working memory is quite stable and robust over time.

On the other hand, other studies have reported finding modest but reliable decrements in discrimination performance over time when evaluating working memory for orientation (Magnussen, Landrø, & Johnsen, 1985), vernier offset (Fahle & Harris, 1992), contrast (Lee & Harris, 1996; Magnussen, Greenlee, & Thomas, 1996), and color (Nilsson & Nelson, 1981). These conflicting findings across studies present a challenge for understanding the processes by which visual representations are sustained over time in working memory.

Although these previous studies relied on established psychophysical procedures, multiple factors complicate their interpretation. Some studies probed memory with a restricted range of test stimuli, for example by testing the discrimination of spatial frequencies centered about a particular reference frequency (with modest jitter applied to the standard reference). As a consequence, delayed discrimination thresholds, in what was deemed a sensory memory task, may have benefited from learning of a "standard" or "criterion" for making these forced-choice discriminations. Such effects of criterion learning, which can last over many trials or even over days, have been shown to contribute to delayed discrimination performance for visual properties such as spatial frequency (Lages & Treisman, 1998; Lages & Paul, 2006) and line separation (Morgan, Watamaniuk, & McKee, 2000). Thus, criterion learning could in some cases explain lossless working memory.

Most previous attempts to characterize the time course of visual working memory suffer from an even greater limitation, namely the inability to distinguish between whether the precision of visual working memory for an item gradually decays over time, or whether on a subset of trials, memory for an item is abruptly lost at some point during the maintenance period—due to lapses of attention, loss of motivation, or any other factor. For example, if participants have a small but nonzero probability of experiencing an attentional lapse within any fixed time interval, trials with longer delay intervals will lead to more frequent lapses. This would consequently lead to elevated estimates of discrimination thresholds at longer delays. Thus, attentional lapses would be expected to impair working memory performance as a function of temporal delay, in both delayed discrimination tasks (Magnussen et al., 1985; Vogels & Orban, 1986) and change detection tasks (Phillips, 1974). Most previous studies of working memory have not adequately distinguished between these two possibilities, namely whether items in working memory gradually degrade over time because of accumulating noise or whether they are more likely to be dropped from memory over longer delay periods.

A study by Zhang and Luck (2009) directly addressed the latter issues by applying a mixture model analysis that assumes separate parameters to estimate the precision of working memory for successfully maintained items and the likelihood of complete memory failure. On each trial, participants were presented with three randomly chosen colors to maintain over a variable delay period, after which they were cued to report an item from memory by indicating the remembered color on a continuous color wheel. This method of continuous report allowed for the construction of error distributions, which could then be fitted by assuming a mixture model, comprised of a von Mises distribution to estimate the precision of responses for successfully maintained items, and a uniform distribution to estimate the frequency of trials involving random guessing. The authors found that guessing-related responses became more prevalent at longer delays but that the precision of visual working memory remained stable over time. These results were taken to suggest that visual information does not gradually degrade over time, but rather, is abruptly lost from working memory in a probabilistic all-or-none manner.

This study addressed major shortcomings of previous work, but how should one interpret a reported null effect of temporal delay on memory precision? One possibility is that working memory precision indeed remains highly stable over time, but another possibility is that considerable statistical power is needed to detect the presence of these effects (Skottun, 2004). A substantial amount of internal noise may need to accumulate over an extended period to result in detectable changes in memory precision, especially if the magnitude of this independent noise source is small in comparison to the noise level associated with the initially encoded item in working memory at the start of the delay. Such a prediction arises from the simple fact that the pooled variance of two independent noise sources will reflect their sum. For example, if the standard deviation associated with memory for an item was equivalent to 20° at Time 1, and an additional 10° of independent noise accumulated between Time Points 1 and 2, then memory performance at Time 2 would result in a standard deviation of 22.36° (i.e., the square root of $20^2 + 10^2$) – a value only modestly greater than that observed at Time 1. This example illustrates how a modest level of independent noise, added to an existing noisy representation, may prove challenging to detect. If the same amount of independent noise were added to a memory representation that was initially more precise, with say a standard deviation of 10° at Time 1, then memory performance at Time 2 would be degraded by a greater proportional extent (with a variance of 200° or standard deviation of 14.14°).

Thus far, we have considered the disruptive effects of attentional lapses and the degree of precision loss that would be expected to occur due to noise accumulation. However, another potential cause of delay effects could be interference in the time domain. The "temporal interference model" posits that representations in memory will become less distinct as they recede into the past, losing temporal distinctiveness. Presumably, this is because of logarithmic psychological compression of time (Brown, Neath, & Chater, 2007). If the delay duration were to vary across trials in a study while the interval between trials remains fixed, then temporal distinctiveness will systematically covary with delay duration. Consider a trial with a long memory delay period: as the sample item recedes into the past it will become less distinguishable from the sample item that appeared on the trial before it. A recent study of visual working memory has found some evidence to support the temporal distinctiveness hypothesis (Souza & Oberauer, 2015). Thus, while worse performance at longer delays could be due to memory decay, decay may be conflated with the loss of temporal distinctiveness, which could lead to greater competition for retrieval between items remembered in the past.

The primary goal of our study was to provide a rigorous test of whether information maintained in visual working memory is subject to gradual loss of precision over time. That guess rates increase at longer delays has been fairly well established (Park, Sy, Hong, & Tong, in press; Zhang & Luck, 2009), and if one assumes that lapses of attention can lead to complete memory failure, an increase in guess rate is indeed expected as a function of delay duration. However, such catastrophic lapses might occur in parallel with a more gradual process of decay that is much harder to detect (see also supplementary Table 1). To enhance sensitivity for detecting potential effects of temporal delay on precision, we evaluated visual working memory for single items, which is known to lead to superior encoding precision and lower rates of memory failure when compared with memory for multiple items (Bays & Husain, 2008; Pratte, Park, Rademaker, & Tong, 2017; Zhang & Luck, 2008).

An important consideration is that one's ability to detect any inherent change in memory precision, based on a mixture model analysis, will depend greatly on the baseline level of memory precision as well as the overall rate of guessing. Statistical power to detect small changes in memory precision will be best when baseline precision is high (i.e., low standard deviation) and guess rates are low. To gain a quantitative appreciation of this, we performed a simulation analysis, generating data for in-memory responses with an initial standard deviation of 10°, 20° or 30° for Time 1, and a proportional increase in standard deviation of 10% or 20% for Time 2. The probability of random-guess responses was varied from 0 to 0.5. We fit the mixture model to the simulated data (100 trials per delay condition), which we repeated for 1000 simulations. Figures 1A and 1B show estimated power for detecting a statistical difference following 10% and 20% increases in standard deviation, respectively, assuming a sample size of $N = 12$ and two-tailed alpha of 0.05. Not only is power greatest when baseline standard deviation is low and guessing is at a minimum, the impact of baseline precision becomes magnified as guess rates increase much above 0. In Figure 1C, we plot the estimated parameters for standard deviation at three levels of guessing (pU) for the data simulated in Figure 1B; the amount of overlap between error bars for Time Points 1 and 2 (T1 and T2) relates to the difficulty/ease with which the impact of delay duration can be detected. This power analysis clarifies our motivation for testing working memory for single items, namely, to ensure that guess rates and baseline standard deviation will be as low as possible. This will improve our ability to detect any change in standard deviation, should memory precision truly decline as a function of delay duration.

We investigated working memory for basic visual features of orientation and color (Experiments 1A and 2, respectively), and performed a control study of memory for orientation to address potential concerns regarding temporal distinctiveness (Experiment 1B). We also characterized the effects of delay duration on visual working memory for complex stimuli, using a continuous face stimulus set (Lorenc, Pratte, Angeloni, & Tong, 2014) to determine the generality of our findings. Across all three types of stimuli, we observed highly consistent effects of declining memory precision as a function of temporal delay, as well as modest increases in the frequency of random-guess responses that reached statistical significance in some cases. These effects were further confirmed by a model comparison analysis that evaluated the likelihood that participants' working memory performance arose from a loss of memory precision over time, as compared with other types of loss such as an increased frequency of complete forgetting. Our results provide compelling evidence that information represented in visual working memory accumulates noise over time, leading to a gradual but inevitable loss of visual precision.

Experiment 1A

In Experiment 1, we evaluated memory fidelity for the orientation of a single grating across a range of delay durations. Orientation is a fundamental visual feature that is prominently represented in early visual areas, and varies in a continuous manner in a circular feature space. A major advantage of testing visual working memory in a full circular feature space is that effects of expectation or criterion learning (Lages & Treisman, 1998) can be eliminated, as any possible feature value within the space can appear as the sample to be remembered. These properties readily allow for evaluation of the precision with which specific orientation values can be maintained over time, and are commonly tested in studies of visual working memory (Luck & Vogel, 1997; Fougnie, Asplund, & Marois, 2010; Fougnie, & Alvarez, 2011; van den Berg, Shin, Chou, George, & Ma, 2012; Rademaker, Tredway, & Tong, 2012; Rademaker, Bloem, De Weerd, & Sack, 2015; Wilken & Ma, 2004). We applied a mixture-model analysis which assumed that successful maintenance would lead to modest errors centered around the feature value of the sample stimulus, while occasional trials involving complete memory failure would be well-described by a uniform-guessing distribution to account for the probability of memory failure in each delay condition (Zhang & Luck, 2008; Pratte et al., 2017).

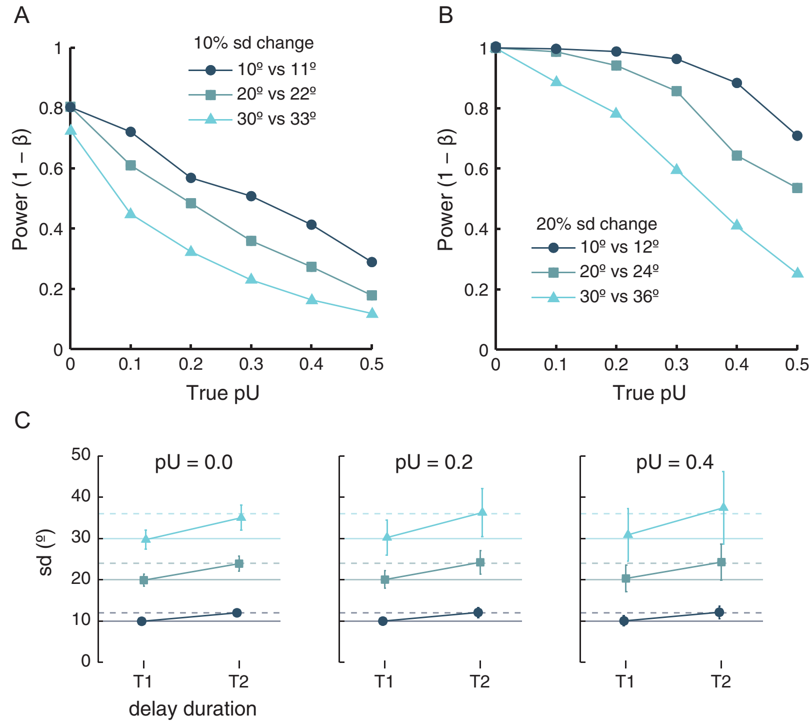

Figure 1. Challenges of detecting gradual decay with mixture model analysis. A simulation analysis was performed assuming a baseline standard deviation of 10°, 20°, or 30° at Time 1 and a proportional increase in standard deviation of 10% (Panel A) or 20% (Panel B) by Time 2. Statistical power to detect a change in standard deviation is optimal when baseline standard deviation is low and guess rates approach 0. Power declines as a function of increasing guess rate, particularly when baseline standard deviation is moderate or high. Panel C: Estimated standard deviation values for the simulation in 1B, plotted for guess rates of 0, 0.2 and 0.4. While the mixture model captures the true standard deviation quite well, these estimates become more variable with larger baseline values of standard deviation and at higher rates of guessing. Error bars show ±1 standard deviation. Note that no variability in effect size was introduced in this analysis; expected power would be lower if the magnitude of change in standard deviation varied across participants. Transparent solid and dashed lines in Panel C represent true standard deviation at Time Points 1 and 2, respectively. Please see the online article for the color version of this figure.

In Experiment 1A, we tested temporal delays of 1, 3, 6, and 12 s using a mixed trial design with a fixed intertrial interval. This study found that visual precision progressively declined as a function of delay duration.

Method

Participants. Twelve healthy volunteers participated at Maastricht University under the approval of the standing ethical committee of the local Psychology and Neuroscience department. All participants reported normal or corrected-to-normal vision, and provided written informed consent. With the exception of one of the authors (RR), participants were naïve to the purpose of the study and received monetary reimbursement for their time. Participants' ages were between 21 and 36 (8 female).

Stimuli. Participants viewed the stimuli in a dark room on a luminance linearized CRT monitor with 1,280 × 1,024 resolution with 60 Hz refresh rate. Visual stimuli were generated using MATLAB and the Psychophysics toolbox (Brainard, 1997; Pelli, 1997). Participants were seated at a viewing distance of 57 cm, and were instructed to maintain fixation throughout, aided by a chinrest and a central fixation bull's eye (0.5° diameter). The orientation stimuli consisted of linear sine-wave gratings (3° diameter; spatial frequency 2 c/°; phase randomized; 20% Michelson contrast with an added 10% contrast jitter) that were centrally presented on a uniform gray background with the same mean luminance (40.8 cd/m^2) The probe stimulus consisted of two line segments (each 0.025° wide and 0.125° long, separated by a 3° gap) indicating a single orientation, and the white bull's eye (0.5° diameter) at fixation. The line could be rotated around the bull's eye using the mouse, allowing participants to replicate the memorized orientation.

Procedure. Participants were presented with a randomly chosen orientation (between 0° and 180° in 1° steps) for 200ms on each trial, followed by a randomly intermixed delay condition of 1, 3, 6, or 12 s (Figure 2A, top row). After the delay participants were presented with a probe in the form of an interrupted line, with an initially random orientation. Participants used a mouse to rotate the line to report the orientation in memory as precisely as possible. Once satisfied with their response, participants clicked the mouse to continue to the next

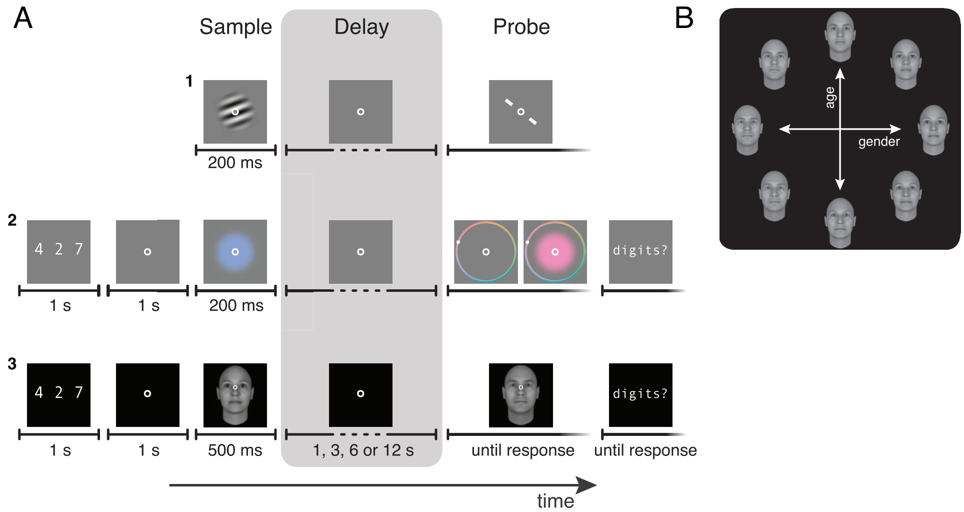

Figure 2. Trial sequence and stimuli. Panel A: In Experiment 1A (top row) participants viewed a randomly oriented grating for 200 ms at the start of each trial, and remembered it as precisely as possible for a duration of 1, 3, 6, or 12 s. After the delay, participants were presented with a randomly oriented probe stimulus, which they could rotate by using the computer mouse to report the orientation from memory. In Experiment 2 (middle row), participants first viewed three digits that they repeated aloud throughout the trial to induce articulatory suppression. After a brief delay, they were shown a randomly colored circular patch for 200 ms, and had to maintain this color information for 1, 3, 6, or 12 s. After the delay, color memory was probed by showing participants a color wheel. A white response circle was shown on the wheel, and once participants started turning the response knob a color patch appeared centrally. Participants could move the white response circle along the color wheel (simultaneously changing the color of the central patch) by turning the response knob until they were satisfied with their response. Finally, participants had to report the three rehearsed digits using a computer keyboard. Experiment 3 (bottom row) evaluated visual working memory for 3D rendered face images. The probe appeared centrally, at the same location as the stimulus, and turning the knob made the face morph through face space to arrive at the desired response. Panel B: Examples of face stimuli that comprised a continuous face space, generated via rendering of 3D faces that varied according to age and gender. For illustrative purposes, the dimension of the stimuli above are not to scale—see the text for actual sizes. Please see the online article for the color version of this figure.

trial 1 s later. For each participant, a total of 100 trials per delay condition were collected.

Analysis. For each delay condition, we first calculated the circular variance (V) – a descriptive statistic measuring the overall dispersion of responses (Mardia & Jupp, 2000). The circular variance was calculated with respect to the true value rather than the mean, as follows:

$$V = \frac{1}{n}\sum_{i=1}^{n}\{1 - \cos(\theta_i - \alpha)\}$$

where α is the true value of the stimulus and θ is the reported value. We also calculated the average response time for each condition of interest. Next, a distribution of response errors was obtained by calculating the difference between sample orientation and the reported orientation. Relevant characteristics from these response error distributions were estimated by fitting a 'mixture-model,' previously proposed to describe various aspects of the working memory system (Zhang & Luck, 2008). The mixture-model assumes that a response error distribution can be described by a mixture of two distributions: A von Mises distribution (i.e., the circular analog of the Gaussian distribution) when the probed item was stored in memory, and a uniform distribution when the item was lost and a random response was made. The von Mises distribution is represented by two parameters, the distribution mean (μ) indicating any systematic shift of the distribution with respect to the correct response, and the standard deviation of the distribution, which is believed to be inversely proportional to memory precision. For these analyses, we assumed that the mean of the distribution was centered at the value of the true stimulus to be remembered. The uniform distribution is represented by one parameter, *pU*, which determines the height of the uniform distribution, indicative of the probability of memory failure.

Data analyses were performed in MATLAB using custom functions, as well as circular data analysis tools provided by the Bays lab (Bays, Catalao, & Husain, 2009) and the Circular Statistics Toolbox (Berens, 2009). We used a maximum likelihood method to estimate the standard deviation and *pU* parameters for each delay condition from each participant. A repeated-measures analysis of variance was then performed on the estimated parameters, with the delay interval as the within-subjects factor. Post hoc *t* tests were uncorrected to provide a more sensitive measure of where differences found with analyses of variance (ANOVAs) might arise from.

Power analysis. In Experiment 1A, our goal was to test for loss of precision in visual memory over temporal delay. As mentioned in the introduction, it may be challenging to detect a modest level of independent noise added to an existing noisy representation, with expected effect sizes in the 2°–4° range (assuming 10° of independent noise, and depending on the noisiness of the existing representation). To be sensitive to such effect sizes, we decided a priori to Test 12 participants for this first experiment, based on our experience of testing similar numbers of participants in previous studies (Lorenc et al., 2014; Pratte et al., 2017; Rademaker et al., 2012, 2015). Next, we calculated the statistical power of detecting a significant increase in the standard deviation, given our sample size and the empirically observed effect size.

Utilizing all the data from Experiment 1A, a slope analysis of the change in the standard deviation parameter as a function of delay revealed a statistically significant slope of + 0.17° *sd*/s (SEM = 0.04° *sd*/s). For more on this analysis see our section on "Analysis of Rate of Precision Loss Across Experiments". Power was calculated for a 1-sample t test against the null hypothesis of slope = 0, and yielded a Cohen's d = 1.345. Assuming similar rates of change in standard deviation for colors and faces, we considered a sample size of ~12 participants for each of the following experiments to provide suitable power.

Results

Results for Experiment 1A are shown in Figure 3. The error histograms (Figure 3A) suggest that reports became less accurate as a single orientation had to be retained over longer delays. To quantify the overall magnitude of memory error, we first calculated the circular variance (V) of response errors in each delay condition (Figure 3B, left most panel, in black, left *y*-axis). Note that either a decline in memory precision or an increase in the proportion of random guesses will lead to an increase in the

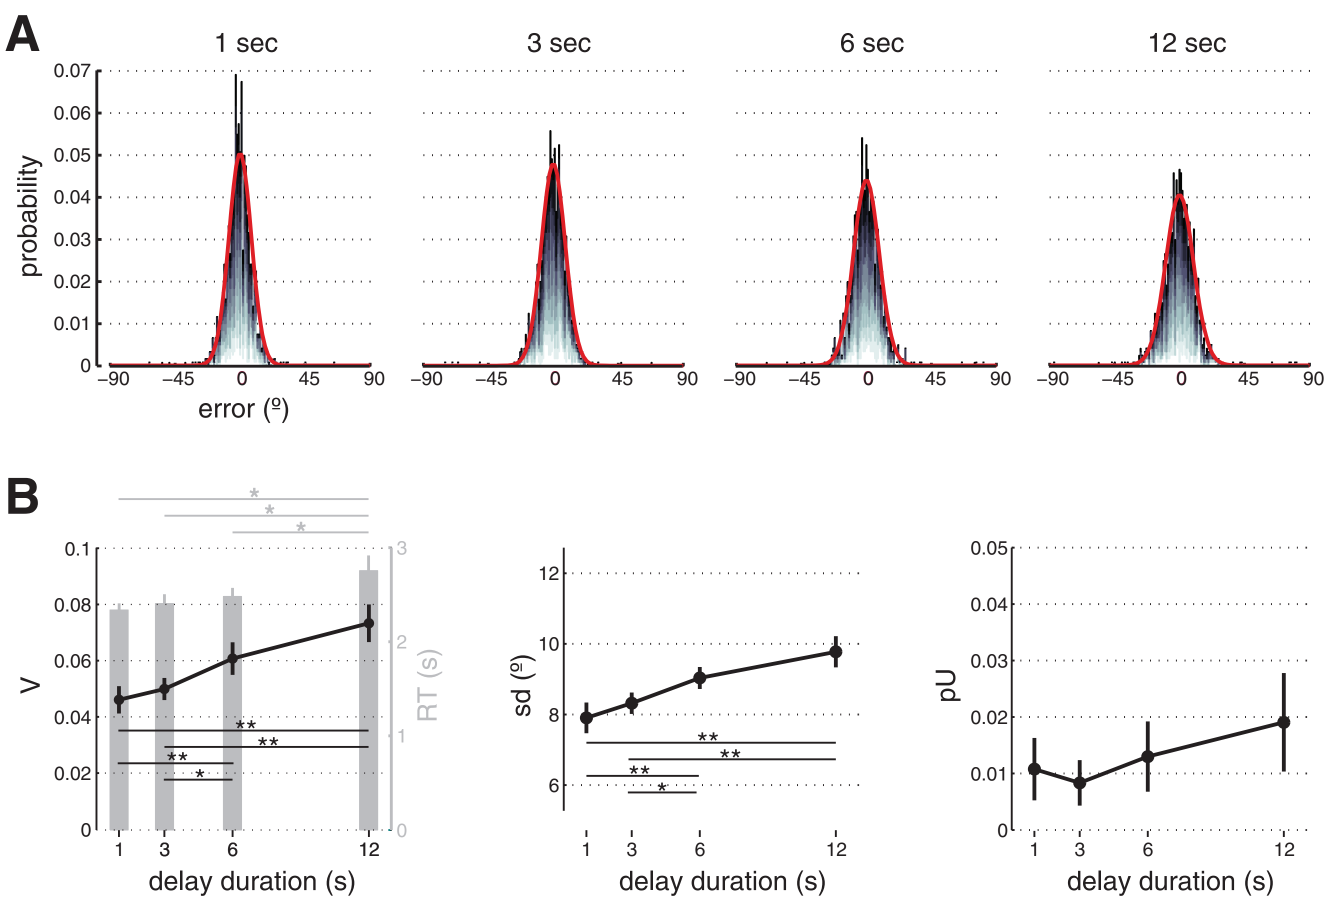

Figure 3. Memory performance for orientation following delays of 1, 3, 6, and 12 s. Panel A: Distribution of errors in working memory responses for each delay duration. Responses became more variable following longer durations. Response frequencies are depicted in gray, and the mixture-model fit across all participants is overlaid in red. Panel B: Left panel shows the circular variance of responses (in black, left *y*-axis) and mean response times (in gray, right *y*-axis), and as a function of temporal delay. The middle panel shows mixture-model estimates of an increase in standard deviation (mixture-model fit), while the right panel shows no change in the probability of uniform responses (*pU*) as a function of delay. Note that the response error histograms (in Panel A) and the standard deviation parameter plot (in Panel B, middle) are based on the orientation space of 0–180°. Error bars indicate ±1 *SEM*. Please see the online article for the color version of this figure.

circular variance of response errors, so for our purposes this metric will serve as a composite measure.

A repeated-measures ANOVA indicated that the circular variance increased with longer retention durations, $F_{(3,33)} = 8.62$; $p < .001$. Paired *t* tests (uncorrected) showed that this difference was significant between delays of 1 and 6 s, $t_{(11)} = 3.9$, $p = .003$; 1 and 12 s, $t_{(11)} = 3.44$, $p = .006$; 3 and 6 s, $t_{(11)} = 2.73$, $p = .02$; and 3 and 12 s, $t_{(11)} = 3.14$, $p = .009$. The average time that participants took to respond during the various delay conditions is also shown in the leftmost panel of Figure 3B (in gray, right *y*-axis). Response times were slower following longer delay durations as indicated by a repeated-measures ANOVA, $F_{(3,33)} = 5.128$; $p = .005$. Particularly, response times were longer when the delay was 12 s, compared with when the delay was 1 s, $t_{(11)} = 2.7$; $p = .02$; 3 s, $t_{(11)} = 2.2$; $p = .05$; or 6 s, $t_{(11)} = 2.42$; $p = .03$.

Next, we used a mixture-model to decompose the error distributions into the standard deviation of response errors for successfully maintained items and the probability of uniform guessing-related responses (*pU*). The middle panel of Figure 3B shows the effect of delay on the precision of visual working memory, where increases in standard deviation over time indicate a gradual loss of precision. A repeated-measures ANOVA confirmed that standard deviation became larger at longer delay durations, $F_{(3,33)} = 11.86$; $p < .001$, with a highly significant linear trend, $F_{(1,11)} = 24.34$; $p < .001$. Paired *t* tests demonstrated significant differences between delays of 1 and 6 s, $t_{(11)} = 4.03$; $p = .002$; 1 and 12 s, $t_{(11)} = 4.43$; $p = .001$; 3 and 6 s, $t_{(11)} = 2.34$; $p = .03$; and 3 and 12 s, $t_{(11)} = 4.2$; $p = .002$.

Although the mean proportion of guessing responses appeared to increase slightly as a function of delay duration (Figure 3B, right), estimated rates of guessing were very low, ranging from just 1–2% of trials across the different delay conditions. These results indicate that participants remained focused on the task on the vast majority of trials. A repeated-measures ANOVA showed no reliable difference between conditions; repeated-measures ANOVA, $F_{(3,33)} = 1.12$; $p = .354$, nor was there a significant linear trend that might imply increased guessing with longer delay durations, $F_{(1,11)} = 1.39$; $p = .263$.

Experiment 1B

In Experiment 1B, we performed a replication study to control for effects of temporal distinctiveness (cf. Souza & Oberauer, 2015). Although it might have been preferable to compare extremely short and long durations in this study to maximize sensitivity, we focused on delay durations of 1 s and 4 s, as longer durations would have required prohibitively long intertrial intervals. Even with this 4-s period for the long duration condition, an intertrial interval of up to 28 s was needed to maintain a constant level of distinctiveness across trials.

Method

Participants. Thirteen participants were tested at the University of California San Diego, and ethical approval was granted by the local institutional review board (IRB). All participants provided written informed consent and had normal or corrected-to-normal vision. Only 10 participants completed both days of testing and have been included in the analyses. All participants were naïve to the purpose of the study and received monetary reimbursement or optional course credit (with exception of one of the authors, RR). Participants' ages were between 19 and 34 (5 female).

Stimuli. Were identical to those used in Experiment 1A with the following exceptions: We used a monitor with 1600 × 1200 resolution and 120 Hz refresh rate; participants were seated at a distance of 50 cm from the screen, and the uniform gray background had a luminance of 51.56 cd/m^2.

Procedure. Experiment 1B was similar to Experiment 1A with the following exceptions: Each trial started with a 200ms beep (600 Hz) immediately before the upcoming grating (shown for 100 ms) to alert participants. The grating was followed either by a short (1 s) or long (4 s) delay, after which participants had 3 s to report the orientation with the mouse probe. Notably, short and long delay durations were randomly interleaved, while the temporal distinctiveness on each trial was kept constant. This was done by calculating the intertrial interval (ITI) dynamically from trial to trial, based on the participant's response time: We adapted the definition of temporal distinctiveness from Souza and Oberauer (2015), calculating distinctiveness as the delay duration on a given trial, divided by the time from the offset of the sample display on the *previous trial* until the onset of the test display on the current trial. Using this definition, we adopted a distinctiveness value of 0.1205 throughout all trials, and calculated the ITI between the upcoming trial (*n*) and the previous trial (*n* – 1) as follows:

$$iti_{(n-1:n)} = \frac{(delay_n - distinctiveness \times (delay_{n-1} + RT + Sample + delay_n))}{distinctiveness}$$

Where *reaction time (RT)* stands for the response time on trial *n* – 1, and *Sample* stands for the duration of the sample stimulus (100 ms). This procedure results in four different trial types: "short–short", "long–long", "short–long", and "long–short". As can be seen in Figure 4, especially a trial with a long delay (4 s) that is preceded by a trial with a short delay (1 s) requires a considerably long ITI between the two trials to keep distinctiveness constant. Also in Figure 4, the response time summed with the ITI is constant within each trial type. If participants did not respond within 3 s, a low beep was played (300 Hz) and they automatically progressed to the next trial. Participants performed 10–50 practice trials before starting the main task to get used to the restricted response time. Analyses were performed on trials in which the participant responded within the required time window, and the proportion of excluded trials was low (4.5%). Participants completed 4 blocks of 100 trials over the course of 2 days, resulting in 200 trials per condition.

Results

In Experiment 1B, we tested delay durations of 1 and 4 s while adjusting the duration of intertrial interval prior to each trial to maintain a constant level of temporal distinctiveness. The effect of delay duration in this experiment was highly consistent with our findings in Experiment 1A.

As can be seen in Figure 5, we observed a statistically significant increase in standard deviation between delays of 1 and 4 s, $t_{(9)} = 3$; $p = .015$, indicating a loss of memory precision even when the distinctiveness of the sample orientation is held constant. Both the values of standard deviation and the magnitude of its increase between 1 and 4 s (mean values of 7.32° and 8.44°

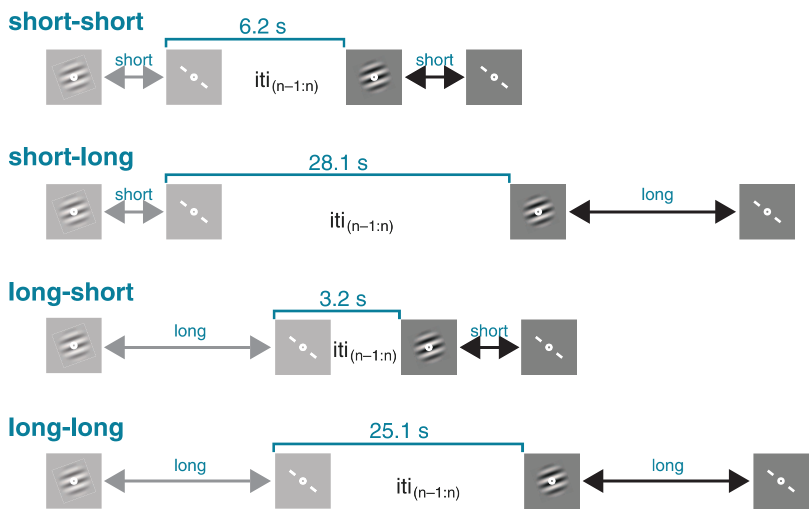

Figure 4. Constant temporal distinctiveness in Experiment 1B. Two delay durations, 1 and 4 s, were randomly interleaved within single blocks of trials, resulting in four trial types: a trial *n* with a short delay preceded by a trial n–1 with a short delay (short-short), a long delay *n* preceded by a short delay n–1 (short-long), a short delay *n* preceded by a long delay n–1 (long-short), and a long delay *n* preceded by a long delay n–1 (long-long). Trials of all four types had equal distinctiveness due to dynamically calculating the intertrial interval from trial-to-trial. Please see the online article for the color version of this figure.

respectively, difference of 1.12°) are consistent with the performance observed in Experiment 1A.

Similar to Experiment 1A, there was no significant effect of delay duration on the probability of guessing, $t_{(9)} = 0.95$; $p = .367$. We observed a marginally significant increase in the circular variance *V*, $t_{(9)} = 2.1$; $p = .064$, at longer delays, presumably driven by the increase in standard deviation that was found. Response times, though capped by the 3-s maximum, were significantly slower at longer delays, $t_{(9)} = 3.57$; $p = .006$. The number of trials on which participants did not respond within the allowed response window did not differ between 1- and 4-s delay conditions, $t_{(9)} = 0.16$; $p = .875$; 89 and 91 total trials respectively across all participants. The slower response times at longer delays, which accompanied the observed increase in standard deviation, suggested that participants found the long delay condition to be more challenging. Taken together, the results of Experiments 1A and 1B demonstrate a gradual loss of precision for orientation information held in working memory, even when temporal distinctiveness is held constant.

Experiment 2

In Experiment 2, we evaluated visual working memory for color over prolonged delays, to determine whether memory for a single color item might also reveal evidence of a gradual loss of precision over time. Our color displays were appropriately calibrated using a spectrophotometer to measure spectra of the red, green and blue phosphors of our CRT monitor, so we could

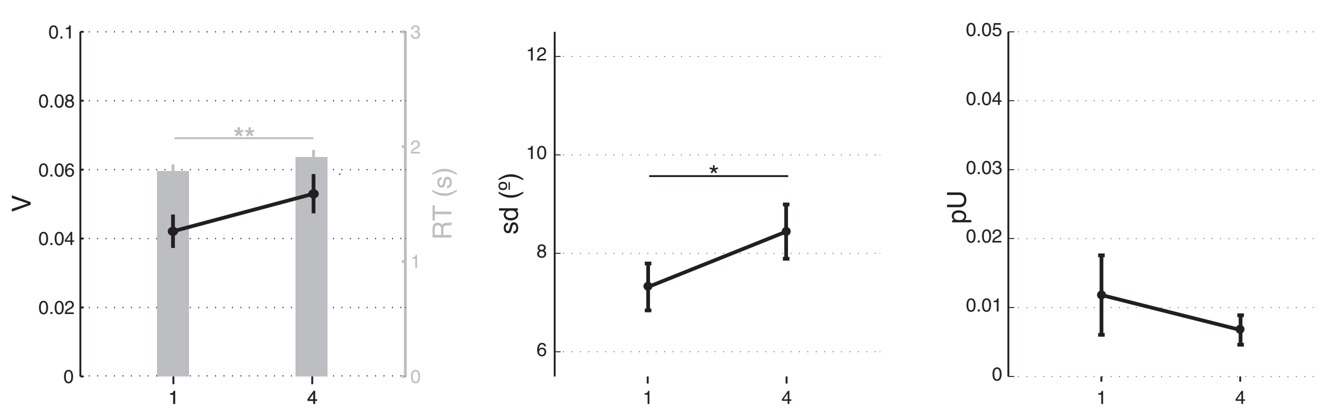

Figure 5. Memory performance for orientation when temporal distinctiveness is held constant across trials. Circular variance, response times, and mixture-model standard deviation all increase as the delay is lengthened from 1 to 4 s, collectively implying less precise memory representations over time. The probability of guessing responses is unaffected by the duration of the memory delay. Error bars indicate ±1 *SEM*.

generate appropriate colors in CIE-LAB space. A proper calibration procedure was important to ensure that the rendered colors were isoluminant and adjacent colors were separated by a constant hue angle difference (Bae, Olkkonen, Allred, & Flombaum, 2015). Because colors can be quite well remembered by relying on verbal strategies (e.g., Donkin, Nosofsky, Gold, & Shiffrin, 2015), we had participants perform a concurrent verbal working memory task while maintaining colors in visual working memory.

Method

Participants. Participants were 12 healthy volunteers (9 female) recruited at Vanderbilt University, where the study was performed with approval of the Institutional Review Board of Vanderbilt University. Participants provided their informed consent, had normal or corrected-to-normal visual acuity and color vision, and were uninformed about the purpose of the study. Participants, aged between 19 and 26, completed two 1-hr sessions, and received monetary reimbursement for their participation.

Stimuli. Stimuli were generated using MATLAB and the Psychophysics toolbox (Brainard, 1997; Pelli, 1997) and viewed from a 46 cm distance on a color and luminance-calibrated CRT monitor (1152 × 870 resolution, 75 Hz refresh rate) in a darkened room. For color-calibration, the spectra of the monitor RGB primaries were measured at maximum intensity, using an Ocean Optics USB4000 spectrometer. The gamma function of each channel was measured with a Minolta LS-110 luminance meter. Based on these calibration data, CIE $L^*a^*b^*$ coordinates were converted to device-dependent RGB values, using the monitor white point (CIE (x, y) = 0.3208, 0.3104; luminance = 58.56 cd/m^2) as reference white.

Participants were instructed to maintain fixation, which was supported by a chinrest and a central bull's eye fixation (0.48° diameter, white). The color stimulus consisted of a centrally presented circular color patch (3° diameter), with the color randomly chosen from one of 360 color values evenly distributed along a circle in CIE $L^*a^*b^*$ space (centered at $L^* = 70$, $a^* = 0$, $b^* = 0$, with a radius of 42 units). An example of the stimulus and experimental design is shown in Figure 2A, middle row. Stimuli were presented against an equiluminant, achromatic background, which corresponded to the center to the sampled color space. A Gaussian envelope ($SD = 3°$) was applied to the circular color patches, such that saturation fell smoothly from the center toward the rim. The probe display initially consisted of a small white circle (0.4° diameter) placed on a color wheel (11.6° inner and 12° outer diameter). Participants could move the white circle along the color wheel by turning a knob interface (PowerMate 3.0, Griffin Technology, U.S.A.). Once a response was initiated, a test color patch appeared at fixation, instantaneously reflecting the color value indicated by the white circle. The color wheel was randomly rotated from trial to trial, as was the initial position of the white circle, to avoid any systematic correspondence between spatial location and color.

Procedure. While the grating orientations used in Experiment 1 map naturally onto a circular feature space that is difficult to verbalize, colors lend themselves well to the use of verbal strategies. We therefore introduced an articulatory suppression component to this experiment. At the start of each trial, participants were presented with three randomly selected digits for 1 s, which they were instructed to repeat aloud throughout the trial. After a 1-s interstimulus interval, the visual component of the trial followed, during which participants first viewed a randomly chosen sample color (between 0° and 360°) for 200 ms, and then had to maintain this information over a randomly chosen delay of 1, 3, 6, or 12-s before giving an unspeeded response. The color wheel with its white marker appeared after the delay, cuing the participant to report the previous color from memory. Upon the first movement of the knob, a color patch appeared centrally, and the participant could dynamically change the color of this patch by rotating the knob. The color wheel helped to the participant to navigate through the color feature space. Once satisfied with their report color, participants pressed the space bar to continue, at which point they used the keyboard to input the three digits they had been repeating aloud until then. The next trial followed 1-s later. A total of 90 trials were obtained in each delay condition, from each participant. Analyses were performed as described for Experiment 1.

Results

Performance on the digit rehearsal task was nearly perfect across all delay conditions (overall accuracy 99.5%, $SE = 0.17\%$), indicating that our participants rehearsed the digits reliably over the delay period. The circular variance (V) of color report errors is shown in the left most panel of Figure 6 (in black, left *y*-axis), and as can be seen, there is a clear trend of increasing magnitudes of memory error following longer delays. A repeated-measures ANOVA confirmed that the circular variance is greater at longer delays, $F_{(3,33)} = 6.88$; $p = .01$. This increase in the circular variance was further explored using uncorrected paired *t* tests, which indicated reliable differences between delays of 1 and 3 s, $t_{(11)} = 2.44$; $p = .03$; 1 and 6 s, $t_{(11)} = 3.48$; $p = .005$; and 1 and 12 s, $t_{(11)} = 3.68$; $p = .004$; as well as between delays of 3 and 12 s—$t_{(11)} = 2.35$; $p = .04$. The leftmost panel of Figure 6 also shows the response times in light gray bars (right *y*-axis). Response times became longer as a function of delay duration, $F_{(3,33)} = 16.38$; $p < .001$; all paired *t* test $p < .008$. Together, these results indicate that the overall accuracy of color responses becomes worse at longer delays.

The middle and right most panel of Figure 6 show the results of the mixture model analysis. With longer retention periods, there was a gradual loss of memory precision, $F_{(3,33)} = 6.05$; $p = .002$, as well as a significant increase in the probability of guessing-related responses, $F_{(3,33)} = 4.83$; $p = .007$. Both showed significant linear trends over the different delay conditions, $F_{(1,11)} = 15.08$; $p = .003$, and $F_{(1,11)} = 8.4$; $p = .014$, respectively. Paired *t* tests indicated a statistically significant difference in precision between the shortest 1-s delay condition and all other delays, 3-s $t_{(11)} = 2.44$, $p = .03$; 6-s $t_{(11)} = 5.3$, $p < .001$; 12-s $t_{(11)} = 3.78$, $p = .003$. Guess responses were more prevalent for delay durations of 12 s compared with 1 s, $t_{(11)} = 2.91$, $p = .014$; and 3, $t_{(11)} = 2.22$, $p = .048$; and for a delay of 6 compared with 1 s, $t_{(11)} = 2.38$; $p = .037$. These results indicate that maintaining a specific color over prolonged delays leads to a gradual loss of memory precision, as well as a greater likelihood of complete memory failure.

Experiment 3

In Experiment 3, our goal was to determine whether a similar loss of visual precision over time would be evident when participants had to retain a complex object over time. To test memory for

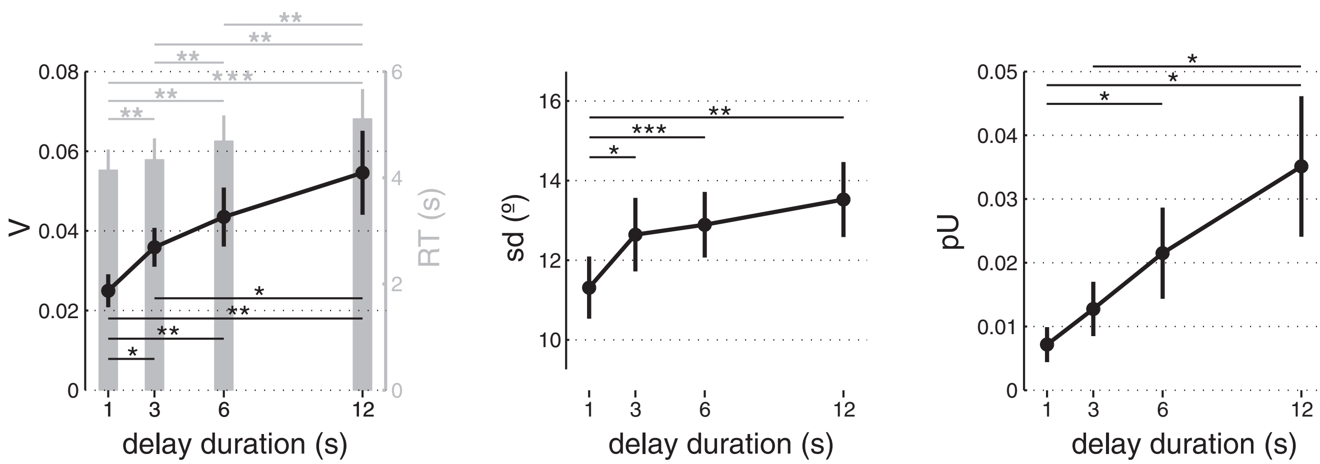

Figure 6. Memory performance for color over various delay durations. Left panel shows an increase in both response time and circular variance for color reports when a single colored patch was maintained in memory for variable delay durations. Mixture-model parameters show an increase in standard deviation (middle panel) as well as an increase in the probability of uniform guess responses (right panel) as a function of delay. Error bars indicate ±1 *SEM*.

complex objects, we developed a set of computer-generated face stimuli (Figure 2B) that continuously varied along the dimensions of gender and age (Lorenc et al., 2014). Unlike previous studies of working memory for complex objects that relied on detection of discrete changes between object stimuli (e.g., Alvarez & Cavanagh, 2004; Awh, Barton, & Vogel, 2007; Banko et al., 2009), our creation of a continuous face space allowed for quantification of the precision of working memory for complex objects, as well as the rate of memory failure (see also Zhang & Luck, 2009).

Method

Participants. Data sets from 13 healthy participants were obtained, but one participant was excluded based on extremely poor overall performance (>2 standard deviation from the group mean) leaving 12 participants included in the analyses. The experiment took place at Vanderbilt University and was approved by the Institutional Review Board of Vanderbilt University. Participants provided written informed consent and were reimbursed monetarily. All participants (ages between 18 and 32; 7 female) had normal or corrected-to-normal vision and were unaware of the goal of the study. Each participant completed a total of 360 trials, distributed across two 1-hr sessions.

Stimuli. Visual stimuli were generated and viewed in the same manner as in Experiment 2, with the exception that one of the monitors was replaced midway through the experiment to one of 1400 × 1050 resolution and 60 Hz refresh rate. Stimuli consisted of gray-scale 3D face images (4.72° by 7.36°) generated with FaceGen Modeler software (Singular Inversions Inc.) as in Lorenc, et al. (2014), and presented against a black background (0.08 cd/m^2 and 1.84 cd/m^2 for the two monitors). Face images were normalized to equate for mean luminance. Eight faces varying along dimensions of age and gender (Figure 2B), forming an octagonal space, were generated first. Next, each pair of neighboring faces were morphed together linearly in varying proportions (10/90, 20/80 . . . 90/10), ultimately resulting in a set of 80 unique faces which we consider to be spaced evenly along a 360° approximately circular 'face space.' This implies that for our analyses, each face in this space is 4.5° apart from its neighbor. The nasal bridge region of each face stimulus was positioned at the center of the screen, with a gray bull's eye fixation (0.08° inner and 0.48° outer diameter) superimposed.

Procedure. Similar to colors, faces can be remembered via verbal strategies, as is custom in day-to-day life. Therefore, a verbal suppression component identical to that of Experiment 2 was used (Figure 2A, bottom row). The visual memory part of each trial started with a 500-ms presentation of a randomly chosen face stimulus, followed by a 1-, 3-, 6-, or 12-s delay. After the delay, a centrally presented probe face morphed through the face-space as participants turned the knob so they could report the face that best matched their memory. Pressing the space bar allowed them to continue to the next trial. As in Experiment 2, each participant completed 90 trials in each delay condition. Analyses were performed as described for Experiment 1.

Results

The performance on the digit task was again highly accurate in all delay conditions (overall accuracy 97.6%, $SE = 0.81\%$), suggesting that all our subjects fulfilled the requirement of articulatory suppression. An analysis of response times indicated that participants required more time following longer delay periods to report a face from memory, $F_{(3,33)} = 15.95$; $p < .001$; all paired t test $p < .03$, Figure 7 left panel. The circular variance of response errors also increased with longer retention durations, $F_{(3,33)} = 8.51$; $p < .001$, and this difference was significant when comparing 1-s with 6-s delays, $t_{(11)} = 2.34$, $p = .039$; 1-s with 12-s delays, $t_{(11)} = 4.62$, $p < .001$; 3-s with 12-s delays, $t_{(11)} = 3.3$, $p = .007$; 6-s with 12-s delays, $t_{(11)} = 3.27$, $p = .007$.

The precision of working memory for faces tended to decline over time, as indicated by an increase in standard deviation at longer delays (Figure 7, middle panel). Although the basic

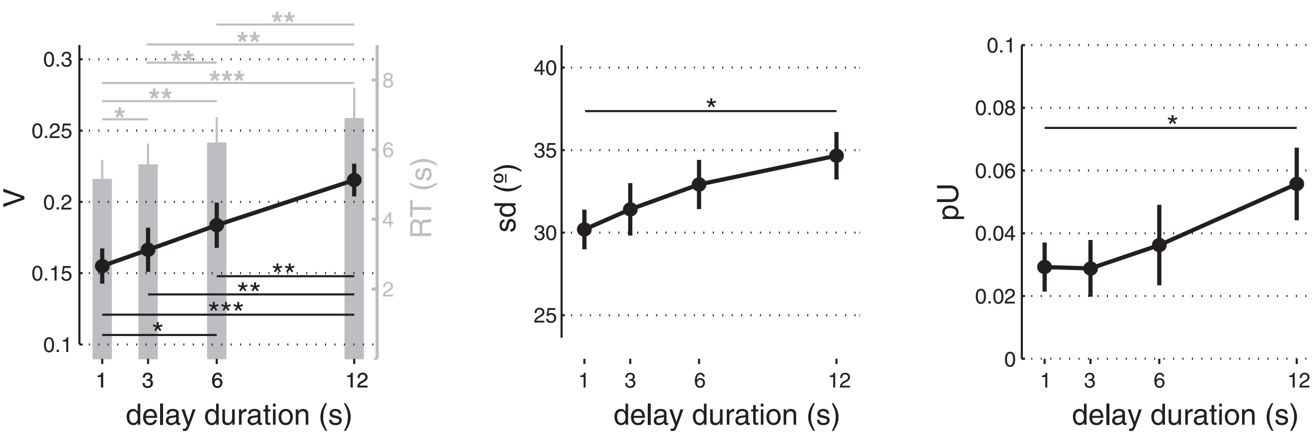

Figure 7. Memory performance for faces over various delay durations. Left panel shows an increase in both response time and circular variance of responses when a single face was maintained over variable delays. The standard deviation showed a trend indicating a decrease in memory precision (middle panel). While the ANOVA revealed no significant change in the probability of uniform guessing responses (right panel), uncorrected paired *t* tests indicated an increase in forgetting between 1- and 12-s delay durations. Error bars indicate ±1 *SEM*.

repeated-measures ANOVA was only marginally significant, $F_{(3,33)} = 2.705$; $p = .061$, we observed a significant linear contrast applied to these data, $F_{(1,11)} = 7.622$; $p = .019$, which was further corroborated by uncorrected paired *t* tests showing that standard deviation increased from 1 to 12 s, $t_{(11)} = 2.68$; $p = .022$. Delay duration did not significantly affect the proportion of guesses (Figure 7, right panel), $F_{(3,33)} = 1.77$; $p = .172$, when evaluated by the ANOVA. However, the idea that participants might guess more often as time wears on was suggested by a statistically significant linear trend, $F_{(1,11)} = 8.866$; $p = .013$, as well as uncorrected paired *t* tests showing more guesses at 12-s compared with 1-s delays, $t_{(11)} = 2.67$; $p = .022$. Overall, the results of Experiment 3 corroborate strongly with those of Experiments 1 and 2, by revealing a gradual loss of memory precision for faces as a function of temporal delay.

Recall errors from all three experiments showed systematic biases, such that error magnitude and sign were nonuniform across the stimulus space (like biases reported in Bae et al., 2015; Pratte et al., 2017). Such biases were not systematically changed or amplified at the different delay conditions (supplementary Figures 1 & 2).

Analysis of Rate of Precision Loss Across Experiments

We performed additional analyses to estimate the rate at which memory precision declined over time for each of the stimulus types across Experiments 1A, 2, and 3 (which have highly similar paradigms). This was done by analyzing the standard deviation values (in °, within each features' respective space) at each time point for individual participants, calculating individual slopes of the best-fitting line to determine the rate of change in standard deviation over time, and then performing a group-level *t* test to determine whether slope values significantly differed from zero. Figure 8 shows the results of this analysis, with the linear fit for individual participants indicated by thin colored lines, and the linear fit averaged across participants indicated by a thick red line. This analysis revealed a statistically significant loss of memory precision in all three experiments (i.e., increase in mixture-model *sd*) with slope values of + 0.17° *sd*/s for orientation (SEM = 0.04° *sd*/s), $t_{(11)} = 4.66$, $p < .001$; +0.17°*sd/*s for color (SEM = 0.05° *sd*/s, $t_{(11)} = 3.28$; $p = .007$; and + 0.40° *sd*/s for faces (SEM = 0.16° *sd*/s), $t_{(11)} = 2.44$; $p = .033$. For comparison, we also fit a square-root temporal decay function to these same data, motivated by the fact the accumulation of independent noise over time should lead to a square-root increase in standard deviation (i.e., a linear increase in variance) over time. The resulting fits of the square-root model (black dashed line, group average) were very similar to the predicted linear fits, as indicated by the high degree of overlap among these fitted curves, and comparable goodness-of-fit values (mean R^2 of 0.655, 0.468, and 0.421 for linear fits of Experiments 1–3, respectively; mean R^2 of 0.659, 0.473, and 0.419 for square root fits of Experiments 1–3, respectively). Because of the high level of variance present at the shortest delay (i.e., the 'intercept' being larger than 0), the amount of bowing that occurs with the square root function is not obvious. Both models provide reasonable quantitative fits of loss of working memory precision over time, and the consistency of these effects is evident across the three stimulus types.

These results were corroborated by an analysis using mixed-effects models (Pinheiro & Bates, 2000). Such models, in addition to the fixed effects of a standard regression model, include random effects that allow for possible interindividual differences in baseline precision or slope. Delay duration was entered into the analysis as a fixed effect, and participant-independent random intercepts and slopes were added incrementally. The likelihood ratio chi-square test was used to determine whether adding the participant-specific intercepts or slopes improved the model.

The model including a fixed effect of delay plus random intercepts for individual participants revealed a significant increase of standard deviation with increasing delay duration in all three experiments, with slope estimates similar to those reported above; orientation: $\beta = +0.17°/s$, $SE = 0.03°$, $t(46) = 6.07$, $p < .001$; color: $\beta = +0.17°/s$, $SE = 0.04°/s$, $t(46) = 3.72$, $p < .001$; face: $\beta = +0.40°/s$, $SE = 0.13°/s$, $t(46) = 2.91$, $p = .005$. Inclusion of

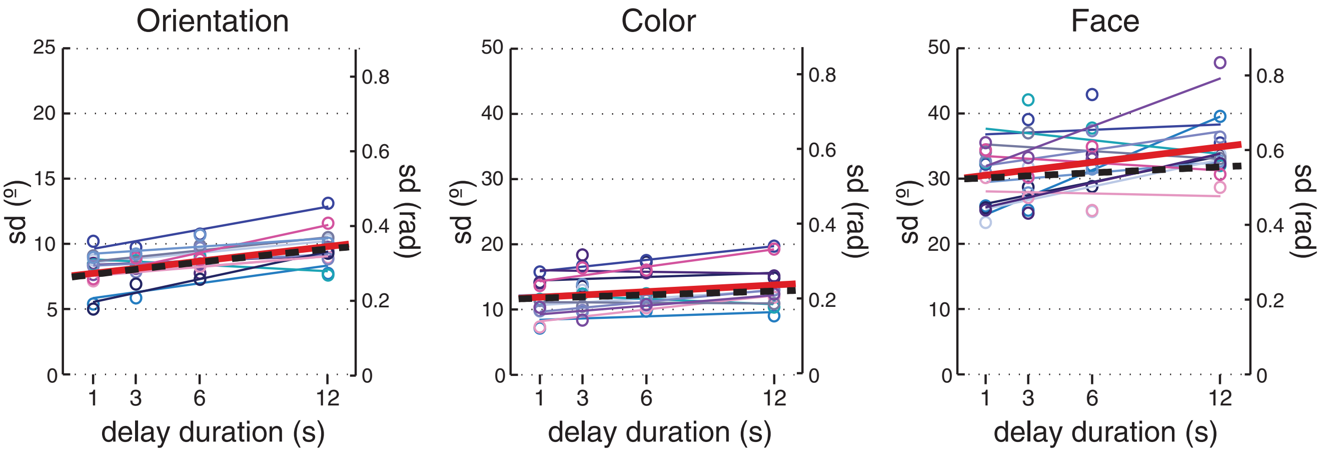

Figure 8. Decay rate of memory precision for Experiments 1–3. Panel A: Plots show individual data points (open circles) and linear fits for individual participants (thin colored lines). Panel B: The average linear fit (thick red line) and the average square-root fit across participants (black dashed curve). Left, middle, and right panels show the mixture model standard deviation of working memory reports for orientation (0–180° space), color (0–360°) and face stimuli (0–360°), respectively. The ordinate axis on the left shows the mixture-model standard deviation of memory errors in degree units for each feature in its respective space, while the right ordinate axis shows the standard deviation of memory errors for all features based on a common circular (0–2π) space in radian units. Please see the online article for the color version of this figure.

the random intercepts component was justified by the likelihood ratio test, orientation: $\chi^2(1) = 20.21$, $p < .001$, color: $\chi^2(1) = 41.91$, $p < .001$; face: $\chi^2(1) = 6.14$, $p = .013$, with estimated standard deviations of the random intercepts (i.e., between-subjects variability in baseline memory precision) of 0.97°, 2.59°, and 2.74°, for orientation, color, and faces, respectively. On the other hand, adding random slopes (alone or along with random intercepts) did not significantly improve goodness of fit for the model—orientation: χ^2s(1) < 1.43, *p*s > 0.232—and was therefore excluded. Thus, the mixed-effects models corroborate our conclusion that memory precision is lost over time for all three stimulus-types tested here.

Model Comparison Analysis of Experiments 1–3

Our statistical analyses of Experiments 1–3 clearly demonstrate reliable changes in working memory performance as a function of delay duration, with consistent positive findings of increasing standard deviation over time. Given that these statistical analyses were performed on parameter estimates obtained by fitting a mixture model to error distribution data, it is worth considering whether a more direct approach to evaluating the data might be possible. One such option is to adopt a model comparison approach, by comparing the quality of fits for different models applied directly to each participant's data.

We constructed a set of four models that specified how memory precision (standard deviation) or guess rate (*pU*) should vary as a function of delay duration. The *Null* model assumes no loss of information over time, requiring that the parameters standard deviation and *pU* are held constant across delay intervals. The *Gradual Decay* model assumes that standard deviation increases as a function of delay duration, while *pU* remains constant. In contrast, the *Sudden Death* model only allows *pU* to increase as a function of duration, but not standard deviation. Finally, the *Hybrid* model allows both standard deviation and *pU* to change, assuming that a memory representation can undergo both gradual deterioration and complete termination.

For all models, the likelihood of observing a given response error (x) in the ***i***th delay condition is specified by:

$$P_i(x) = (1 - pU_i)\mathrm{VM}(0, sd_i) + pU_i/(2\pi),$$

where sd_i and pU_i respectively denote the memory precision and the guess rate in the ith delay condition. The von Mises distribution, VM, is always centered at 0, assuming no response bias with respect to the true stimulus value.

We assume that, as memory precision decays, standard deviation increases at a constant rate over time, following a linear function of retention interval (t):

$$sd_i = sd_1 + (t_i - t_1)sd_{slope},$$

where sd_1 denotes the baseline precision at the shortest delay tested (t_1), and sd_{slope} denotes the rate of standard deviation change per unit time. Likewise, complete forgetting of an item is modeled as a linear increase in *pU* as a function of t:

$$pU_i = pU_1 + (t_i - t_1)pU_{slope},$$

where pU_1 denotes the baseline guess rate at the shortest delay tested (t_1), and pU_{slope} denotes the rate of *pU* change per unit time.

These linear decay and forgetting functions allowed us to fit the data using a smaller number of free parameters (i.e., sd_1, sd_{slope}, pU_1, and pU_{slope}) compared with fitting the mixture model separately to each delay duration ($sd_1{\sim}sd_4$ and $pU_1{\sim}pU_4$). The linear function was chosen mainly for its simplicity, and also because our previous analyses indicated that the linear function described the effect of delay duration on both standard deviation and *pU* reasonably well. As free parameters, sd_{slope} and pU_{slope} were constrained to have positive values only, to ensure that they provided a better fit only when memory performance became worse as a function of delay duration.

The models were fit to each participant's data using maximum likelihood estimation. As a model comparison statistic, we used the Akaike information criterion (AIC; Akaike, 1974), which is calculated as:

$$\text{AIC} = -2 \ln(L) + 2k,$$

where L is the maximum likelihood value of the fitted model, and k is the number of free parameters in the model. Among candidate models, the model that yields the lowest AIC score is selected as the best model. The AIC takes into account both goodness of fit and parsimony of the model, by rewarding high log-likelihoods and penalizing for extra free parameters.

We calculated the AIC score for each model for each participant, and then summed the AIC scores across all participants from each experiment to obtain a total AIC score for each model. The best fitting model (i.e., lowest) AIC score was subtracted from each model's AIC score, indicating how poorly those models did in comparison to the best fitting model. AIC difference scores for each model from all experiments are reported in Table 1.

We found that the gradual decay model provided the overall best fit, yielding the smallest total AIC value in each of the four experiments. In comparison, the sudden death model fared much worse, outperformed by both null and hybrid models in Experiments 1A and 1B, and by the hybrid model in Experiment 2. Thus, when a single item must be maintained in working memory, gradual decay provides the best characterization as determined by AIC.

We also considered which model was selected as the top model most often, based on individual AIC scores. These individual measures, while more variable than group data, do provide information about data trends at the individual level. Table 1 shows the number of individuals that favored each of the four models in parentheses, displayed by experiment. The Gradual Decay model provided the best fit for 7 of 12 participants and 5 of 10 participants in Experiments 1A and B (orientation), 4 of 12 participants in Experiment 2 (color), and 5 of 12 participants in Experiment 3 (face). Although the Hybrid model was the next best performing model at the group level, the Null model was favored next most often among individual participants, and benefitted from having fewer parameters than the more complex models.

We also directly contrasted the Gradual Decay with the Sudden Death model for those participants showing reliable memory loss over time. For this comparison, we considered only those participants for whom the null model was rejected, and compared only two possible models (Gradual Decay and Sudden Death). The Gradual Decay model provided a better fit than the Sudden Death model in 7 of 9 participants in Experiment 1A, 6 of 6 participants in Experiment 1B, and 5 of 8 participants in Experiments 2 and 3.

Overall, the model comparison results indicate that when participants must maintain a single item in working memory, gradual decay provides the best account of how memory performance changes over time. These results strongly corroborate our repeated-measures and slope analyses showing that standard deviation significantly increases as a function of delay duration. However, it would be inappropriate to conclude that these analyses nullify the statistically significant increase in memory failure we observe at longer delays in Experiments 2 and 3. Sudden termination can also contribute to memory loss, and has been found to increase as a function of delay duration in previous studies (Park et al., in press; Zhang & Luck, 2009).

Discussion

The question of whether information in visual working memory is maintained over time in a lossy or lossless manner has been a matter of longstanding debate. Across all four experiments, we found evidence of gradual loss of memory precision over time, even when the frequency of complete memory loss was taken into account through application of a mixture-model analysis. This loss of precision over time was also observed when temporal distinctiveness was held constant, implying that longer delays are a root cause of more variable memory performance. The consequences of temporal delay on memory precision were evident for simple visual features of color and orientation, and extended to memory for complex face stimuli as well. The frequency of complete memory loss did not reliably change as a function of delay duration when orientation was the remembered feature. However, occurrences of complete memory loss were significantly more frequent at longer delays when a single color or face had to be maintained in memory. This conclusion was further supported by model comparison analysis, which revealed that the gradual decay model outperformed models that assume memory loss through sudden death. Our findings provide compelling evidence that items maintained in visual working memory undergo gradual decay due to the accumulation of random noise over time.

Given our highly consistent and positive findings, why have previous studies reported null effects of precision loss over time? Our simulation analyses indicate an important factor to consider, namely that statistical power to detect a change in memory precision will be impeded when guess rates are high and baseline memory precision is low (see Figure 1). In the current study, we

Table 1
ΔAIC Scores

Model	Experiment 1A Orientation	Experiment 1B Orientation	Experiment 2 Color	Experiment 3 Face
Null	65 (3)	45 (4)	27 (4)	18 (4)
Gradual decay	— (7)	— (5)	— (4)	— (5)
Sudden death	73 (2)	54 (0)	16 (2)	9 (3)
Hybrid	14 (0)	14 (1)	2 (2)	15 (0)

Note. AIC = Akaike information criterion. Difference scores are calculated with respect to the best-fitting model for each experiment, which in all cases proved to be the gradual decay model. Dashes indicate the best fitting model.

focused on working memory for single items, which ensured low rates of memory failure (less than 5%) and superior levels of baseline precision than would otherwise be achieved by testing multiple items. In their main experiment, Zhang and Luck (2009) evaluated working memory for 3 color patches across different delays. They found that guess rates increased significantly from 1 to 10 s (from 0.26 to 0.39) while the increase in standard deviation (from 22.9° to 24.4°) did not approach statistical significance. In our simulation analyses (see Figure 1), we observe a sharp decrease in statistical power as guess rates rise upward to 30 or 40%. Also at these high guess rates, an increase in baseline standard deviation (from say 10° to 20°) leads to a considerable drop in power at detecting a proportional increase in standard deviation of 10% or 20%.

Other studies that tested for delay effects with multiple items have likewise reported null effects of precision loss over time, in situations where guess rates were very high (e.g., Souza & Oberauer, 2015), or guess rates were not estimated so their impact could not be assessed (Pertzov, Bays, Joseph, & Husain, 2013). A recent study evaluated working memory for single colors, and further tested the precision with which participants could reconstruct colors based on the verbal labels they provided on previous trials in response to a color stimulus (Donkin et al., 2015). This study did find positive evidence of a decline in the memory precision at longer delays, of a magnitude that appears consistent with the decay reported here. While these researchers ascribed their effect in terms of a shift from visual memory to a verbal-labeling strategy, we controlled for the use of verbal labels in our Experiments 2 and 3 by requiring articulatory suppression when easy-to-label stimuli (colors, faces) were tested. Our findings across diverse stimulus types suggest that gradual loss of precision is a general and widespread phenomenon.

One might ask whether the precision loss that occurs for single items in working memory might somehow represent a special case. Perhaps working memory for multiple items might somehow avoid mechanisms of gradual decay and noise accumulation, albeit for reasons as yet unknown? We consider such an account unlikely, as our lab has obtained direct evidence to refute it. In a separate study, we investigated working memory for multifeature objects, requiring participants to remember the colors and orientations of two colored gratings for delay durations of 1.5s or 5.5s (Park et al., in press). This study revealed significant effects of gradual decay as well as an effect of sudden termination. Estimated standard deviation for color increased from 19.9° to 21.7° across delays, $t_{(23)} = 5.69$, $p < .001$; Cohen's $d = 1.16$ and standard deviation for orientation increased from 14.7° to 17.1° ($t_{(23)} = 4.16$; $p < .001$; Cohen's $d = 0.85$). In addition, we observed a significant increase in guess rates for color from 7% to 12%, $t_{(23)} = 3.52$; $p = .002$; Cohen's $d = 0.72$, while guess rates for orientation (10% and 11%) did not; change over this interval, $t_{(23)} = 0.52$; $p > .250$. The results of this separate study indicate the gradual decay of information is pervasive, and can be reliably detected when multiple objects must be maintained in working memory. This study also supports the finding that longer delays lead to higher rates of memory failure, as has been reported in previous studies.

Additionally, we ensured that the observed memory decay did not merely reflect differences in temporal distinctiveness. Some recent studies have found that differences in temporal distinctiveness, rather than delay duration, negatively affected measures of working memory performance for large arrays of items (Mercer, 2014; Souza & Oberauer, 2015). In such work, temporal distinctiveness has been numerically described as the ratio of the memory delay to the intertrial interval (Souza & Oberauer, 2015). We adapted this temporal distinctiveness ratio, and kept the ratio constant from trial to trial ensuring equal temporal distinctiveness. Moreover, two delay durations (1 and 4 s) were randomly interleaved within the same blocks of trials. Keeping temporal distinctiveness constant while randomly interleaving trials requires much more time than it would to present the delay conditions in a blocked fashion. Nevertheless, blocked presentations have serious drawbacks, an obvious example being the potential of modifying the participants' level of arousal (or boredom) with long blocks of trials. With our experimental design, we ruled out differences in temporal distinctiveness explaining the delay effect, supporting the idea that memories decay over time.

In addition to the changes in the accuracy of working memory over time, we found that response times were considerably slower following longer delays, increasing in an approximately linear fashion for all three types of stimuli. Longer response times may reflect the fact that participants became less sure of the accuracy of visual working memory with the passage of time, which would be expected if their internal representations became noisier (Nilsson & Nelson, 1981). One might even construe working memory retrieval as a decision process based on noisy evidence, in which case having to accumulate information from noisier representations should lead to slower decisions and longer response times (Pearson, Raškevičius, Bays, Pertzov, & Husain, 2014). Note that response times in our experiments were unspeeded (Experiments 1a, 2, and 3) and thus did not directly test decision-making processes. Nevertheless, these results provide another way of demonstrating that participants found it more difficult to make a response at longer delays, and these effects were highly consistent across our experiments, including Experiment 1b where responses were made under time constraints.

The fact that longer retention periods lead to a consistent decrease in precision is consistent with the proposal that working memory representations gradually accumulate noise over time. These results have important implications for neural models of working memory, which must ultimately address how neural noise should be understood and incorporated to appropriately characterize human working memory performance (Bays, 2014, 2015; Brody, Romo, & Kepecs, 2003; Fougnie, Suchow, & Alvarez, 2012; Simmering, Schutte, & Spencer, 2008; Sreenivasan, Curtis, & D'Esposito, 2014; van den Berg et al., 2012; Wang, 2001; Wei, Wang, & Wang, 2012).

Our findings demonstrate systematic changes in memory precision over time. Resource models of working memory assume that working memory has no discrete capacity limit, but that memory precision for a given item declines in a continuous fashion as a function of set size because working memory resources become thinly divided among the items (Ma, Husain, & Bays, 2014). According to this account, one might expect visual working memory for an item to gradually degrade over time, until the representation becomes so noisy that behavioral accuracy appears similar to random guessing (Bays, 2015; Fougnie et al., 2012; van den Berg, et al., 2012). Our data do not directly speak to the issue of whether the guessing-like responses occurring at longer delays reflect severely degraded memory representations or complete termination of the representations. However, we want to point out that the occurrence of such extreme errors tends to increase as a function of time, suggesting that a memory representation is more likely to undergo an abrupt change in noise level as the representation is maintained over longer periods. The resource-based view can be contrasted with slot-based models of visual working memory. Slot models state that 3–4 discrete items can be stored in working memory (Luck & Vogel, 1997). The slots-plus-averaging account (Zhang & Luck, 2008) further suggests that when a single item is remembered, each slot can be utilized to hold a noisy estimate of that item. If more available slots are dedicated to a single item, then averaging across multiple slots will lead to a more precise estimate. In this framework, it is also possible that the likelihood of a slot being dropped increases as a function of time, causing reduced precision of report. An item will eventually become inaccessible when all slot representations of that item have been dropped.

An important and challenging question for future research concerns the relationship between gradual loss of memory precision and the sudden termination of working memory. Our findings suggest that both mechanisms are likely at play. If the mechanisms underlying noise accumulation and complete memory loss are truly independent, one would expect that memory for an item could suddenly fail, perhaps due to an attentional lapse, regardless of its current noise level. Alternatively, noise accumulation and complete loss might reflect two phases of a common process, whereby noise accumulates first, followed by a complete loss once the noise level exceeds some threshold. It will be interesting to see if future investigations of the evolution of visual working memory over time can help to distinguish between these differing theoretical perspectives.

In summary, our study reveals that basic features and complex objects maintained in working memory undergo a gradual loss of precision over time, as well as an increased likelihood of sudden termination. Consistent decline in precision over time indicates that working memory representations are quite stable but are nevertheless subject to the gradual accumulation of internal noise. Further investigation into the nature of visual working memory and how it evolves over time may provide new insights into the underlying mechanisms and limits, and could prove valuable for understanding how working memory operates and dynamically adapts to ongoing cognitive demands.

References

Akaike, H. (1974). A new look at statistical model identification. *IEEE Transactions on Automatic Control, 19,* 716–723.

Alvarez, G. A., & Cavanagh, P. (2004). The capacity of visual short-term memory is set both by visual information load and by number of objects. *Psychological Science, 15,* 106–111. http://dx.doi.org/10.1111/j.0963-7214.2004.01502006.x

Awh, E., Barton, B., & Vogel, E. K. (2007). Visual working memory represents a fixed number of items regardless of complexity. *Psychological Science, 18,* 622–628. http://dx.doi.org/10.1111/j.1467-9280.2007.01949.x

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,* 744–763. http://dx.doi.org/10.1037/xge0000076

Bankó, E. M., Gál, V., & Vidnyánszky, Z. (2009). Flawless visual short-term memory for facial emotional expressions. *Journal of Vision, 9,* 12.1–12.13. http://dx.doi.org/10.1167/9.1.12

Bays, P. M. (2014). Noise in neural populations accounts for errors in working memory. *The Journal of Neuroscience, 34,* 3632–3645. http://dx.doi.org/10.1523/JNEUROSCI.3204-13.2014

Bays, P. M. (2015). Spikes not slots: Noise in neural populations limits working memory. *Trends in Cognitive Sciences, 19,* 431–438. http://dx.doi.org/10.1016/j.tics.2015.06.004

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,* 7.1–7.11.

Bays, P. M., & Husain, M. (2008, August 8). Dynamic shifts of limited working memory resources in human vision. *Science, 321,* 851–854. http://dx.doi.org/10.1126/science.1158023

Bennett, P. J., & Cortese, F. (1996). Masking of spatial frequency in visual memory depends on distal, not retinal, frequency. *Vision Research, 36,* 233–238. http://dx.doi.org/10.1016/0042-6989(95)00085-E

Berens, P. (2009). CircStat: A MATLAB toolbox for circular statistics. *Journal of Statistical Software, 31,* 1–21. http://dx.doi.org/10.18637/jss.v031.i10

Blake, R., Cepeda, N. J., & Hiris, E. (1997). Memory for visual motion. *Journal of Experimental Psychology: Human Perception and Performance, 23,* 353–369. http://dx.doi.org/10.1037/0096-1523.23.2.353

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

Brody, C. D., Romo, R., & Kepecs, A. (2003). Basic mechanisms for graded persistent activity: Discrete attractors, continuous attractors, and dynamic representations. *Current Opinion in Neurobiology, 13,* 204–211. http://dx.doi.org/10.1016/S0959-4388(03)00050-3

Brown, G. D. A., Neath, I., & Chater, N. (2007). A temporal ratio model of memory. *Psychological Review, 114,* 539–576. http://dx.doi.org/10.1037/0033-295X.114.3.539

Donkin, C., Nosofsky, R., Gold, J., & Shiffrin, R. (2015). Verbal labeling, gradual decay, and sudden death in visual short-term memory. *Psychonomic Bulletin & Review, 22,* 170–178. http://dx.doi.org/10.3758/s13423-014-0675-5

Fahle, M., & Harris, J. P. (1992). Visual memory for vernier offsets. *Vision Research, 32,* 1033–1042. http://dx.doi.org/10.1016/0042-6989(92)90004-3

Fougnie, D., & Alvarez, G. A. (2011). Object features fail independently in visual working memory: Evidence for a probabilistic feature-store model. *Journal of Vision, 11,* 3.1–3.12. http://dx.doi.org/10.1167/11.12.3

Fougnie, D., Asplund, C. L., & Marois, R. (2010). What are the units of storage in visual working memory? *Journal of Vision, 10,* 27.1–27.11. http://dx.doi.org/10.1167/10.12.27

Fougnie, D., Suchow, J. W., & Alvarez, G. A. (2012). Variability in the quality of visual working memory. *Nature Communications, 3,* 1229. http://dx.doi.org/10.1038/ncomms2237

Hayhoe, M., Aivar, P., Shrivastavah, A., & Mruczek, R. (2002). Visual short-term memory and motor planning. *Progress in Brain Research, 140,* 349–363. http://dx.doi.org/10.1016/S0079-6123(02)40062-3

Hayhoe, M. M., Shrivastava, A., Mruczek, R., & Pelz, J. B. (2003). Visual memory and motor planning in a natural task. *Journal of Vision, 3,* 49–63. http://dx.doi.org/10.1167/3.1.6

Hollingworth, A., Richard, A. M., & Luck, S. J. (2008). Understanding the function of visual short-term memory: Transsaccadic memory, object correspondence, and gaze correction. *Journal of Experimental Psychology: General, 137,* 163–181. http://dx.doi.org/10.1037/0096-3445.137.1.163

Huang, J., & Sekuler, R. (2010). Distortions in recall from visual memory: Two classes of attractors at work. *Journal of Vision, 10,* 24.1–24.17.

Lages, M., & Paul, A. (2006). Visual long-term memory for spatial frequency? *Psychonomic Bulletin & Review, 13,* 486–492. http://dx.doi.org/10.3758/BF03193874

Lages, M., & Treisman, M. (1998). Spatial frequency discrimination: Visual long-term memory or criterion setting? *Vision Research, 38,* 557–572. http://dx.doi.org/10.1016/S0042-6989(97)88333-2

Lee, B., & Harris, J. (1996). Contrast transfer characteristics of visual short-term memory. *Vision Research, 36,* 2159–2166. http://dx.doi.org/10.1016/0042-6989(95)00271-5

Lorenc, E. S., Pratte, M. S., Angeloni, C. F., & Tong, F. (2014). Expertise for upright faces improves the precision but not the capacity of visual working memory. *Attention, Perception & Psychophysics, 76,* 1975–1984. http://dx.doi.org/10.3758/s13414-014-0653-z

Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature, 390,* 279–281. http://dx.doi.org/10.1038/36846

Ma, W. J., Husain, M., & Bays, P. M. (2014). Changing concepts of working memory. *Nature Neuroscience, 17,* 347–356. http://dx.doi.org/10.1038/nn.3655

Magnussen, S., & Greenlee, M. W. (1992). Retention and disruption of motion information in visual short-term memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18,* 151–156. http://dx.doi.org/10.1037/0278-7393.18.1.151

Magnussen, S., Greenlee, M. W., Asplund, R., & Dyrnes, S. (1991). Stimulus-specific mechanisms of visual short-term memory. *Vision Research, 31*(7–8), 1213–1219. http://dx.doi.org/10.1016/0042-6989(91)90046-8

Magnussen, S., Greenlee, M. W., & Thomas, J. P. (1996). Parallel processing in visual short-term memory. *Journal of Experimental Psychology: Human Perception and Performance, 22,* 202–212. http://dx.doi.org/10.1037/0096-1523.22.1.202

Magnussen, S., Landrø, N. I., & Johnsen, T. (1985). Visual half-field symmetry in orientation perception. *Perception, 14,* 265–273. http://dx.doi.org/10.1068/p140265

Mardia, K. V., & Jupp, P. E. (2000). *Directional statistics*. New York, NY: Wiley.

Mercer, T. (2014). The loss of short-term visual representations over time: Decay or temporal distinctiveness? *Journal of Experimental Psychology: Human Perception and Performance, 40,* 2281–2288. http://dx.doi.org/10.1037/a0038141

Morgan, M. J., Watamaniuk, S. N., & McKee, S. P. (2000). The use of an implicit standard for measuring discrimination thresholds. *Vision Research, 40,* 2341–2349. http://dx.doi.org/10.1016/S0042-6989(00)00093-6

Nilsson, T. H., & Nelson, T. M. (1981). Delayed monochromatic hue matches indicate characteristics of visual memory. *Journal of Experimental Psychology: Human Perception and Performance, 7,* 141–150. http://dx.doi.org/10.1037/0096-1523.7.1.141

Park, Y. E., Sy, J. L., Hong, S. W., & Tong, F. (in press). Reprioritization of features of multi-dimensional objects stored in visual working memory. *Psychological Science*.

Pearson, B., Raškevičius, J., Bays, P. M., Pertzov, Y., & Husain, M. (2014). Working memory retrieval as a decision process. *Journal of Vision, 14,* 2.1–2.15.

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

Pertzov, Y., Bays, P. M., Joseph, S., & Husain, M. (2013). Rapid forgetting prevented by retrospective attention cues. *Journal of Experimental Psychology: Human Perception and Performance, 39,* 1224–1231. http://dx.doi.org/10.1037/a0030947

Phillips, W. A. (1974). On the distinction between sensory storage and short-term visual memory. *Perception & Psychophysics, 16,* 283–290. http://dx.doi.org/10.3758/BF03203943

Pinheiro, J. C., & Bates, D. M. (2000). *Mixed-effects models in S and S-PLUS*. New York, NY: Springer. http://dx.doi.org/10.1007/978-1-4419-0318-1

Pratte, M. S., Park, Y. E., Rademaker, R. L., & Tong, F. (2017). Accounting for stimulus-specific variation in precision reveals a discrete capacity limit in visual working memory. *Journal of Experimental Psychology: Human Perception and Performance, 43,* 6–17. http://dx.doi.org/10.1037/xhp0000302

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,* 1650–1665. http://dx.doi.org/10.1037/xhp0000110

Rademaker, R. L., Tredway, C. H., & Tong, F. (2012). Introspective judgments predict the precision and likelihood of successful maintenance of visual working memory. *Journal of Vision, 12,* 21.1–21.13. http://dx.doi.org/10.1167/12.13.21

Regan, D. (1985). Storage of spatial-frequency information and spatial-frequency discrimination. *Journal of the Optical Society of America, 2,* 619–621. http://dx.doi.org/10.1364/JOSAA.2.000619

Regan, D., & Beverley, K. I. (1985). Postadaptation orientation discrimination. *Optics and Image Science, 2,* 147–155. http://dx.doi.org/10.1364/JOSAA.2.000147

Simmering, V. R., Schutte, A. R., & Spencer, J. P. (2008). Generalizing the dynamic field theory of spatial cognition across real and developmental time scales. *Brain Research, 1202,* 68–86. http://dx.doi.org/10.1016/j.brainres.2007.06.081

Skottun, B. C. (2004). On the use of discrimination to assess memory. *Perception & Psychophysics, 66,* 1202–1205. http://dx.doi.org/10.3758/BF03196846

Souza, A. S., & Oberauer, K. (2015). Time-based forgetting in visual working memory reflects temporal distinctiveness, not decay. *Psychonomic Bulletin & Review, 22,* 156–162. http://dx.doi.org/10.3758/s13423-014-0652-z

Sreenivasan, K. K., Curtis, C. E., & D'Esposito, M. (2014). Revisiting the role of persistent neural activity during working memory. *Trends in Cognitive Sciences, 18,* 82–89. http://dx.doi.org/10.1016/j.tics.2013.12.001

van den Berg, R., Shin, H., Chou, W.-C., George, R., & Ma, W. J. (2012). Variability in encoding precision accounts for visual short-term memory limitations. *Proceedings of the National Academy of Sciences of the United States of America, 109,* 8780–8785. http://dx.doi.org/10.1073/pnas.1117465109

Vogels, R., & Orban, G. A. (1986). Decision processes in visual discrimination of line orientation. *Journal of Experimental Psychology: Human Perception and Performance, 12,* 115–132. http://dx.doi.org/10.1037/0096-1523.12.2.115

Wang, X. J. (2001). Synaptic reverberation underlying mnemonic persistent activity. *Trends in Neurosciences, 24,* 455–463. http://dx.doi.org/10.1016/S0166-2236(00)01868-3

Wei, Z., Wang, X. J., & Wang, D. H. (2012). From distributed resources to limited slots in multiple-item working memory: A spiking network model with normalization. *The Journal of Neuroscience, 32,* 11228–11240. http://dx.doi.org/10.1523/JNEUROSCI.0735-12.2012

Wilken, P., & Ma, W. J. (2004). A detection theory account of change detection. *Journal of Vision, 4,* 1120–1135. http://dx.doi.org/10.1167/4.12.11

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

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

Received January 23, 2017
Revision received July 7, 2017
Accepted August 25, 2017 ■