Journal of Experimental Psychology: General
2017, Vol. 146, No. 11, 1606–1630

0096-3445/17/$12.00 http://dx.doi.org/10.1037/xge0000359

Erasing and Blurring Memories: The Differential Impact of Interference on Separate Aspects of Forgetting

Sol Z. Sun and Celia Fidalgo
University of Toronto

Morgan D. Barense and Andy C. H. Lee
University of Toronto and Rotman Research Institute at Baycrest Hospital, Toronto, Ontario, Canada

Jonathan S. Cant
University of Toronto

Susanne Ferber
University of Toronto and Rotman Research Institute at Baycrest Hospital, Toronto, Ontario, Canada

Interference disrupts information processing across many timescales, from immediate perception to memory over short and long durations. The widely held *similarity assumption* states that as similarity between interfering information and memory contents increases, so too does the degree of impairment. However, information is lost from memory in different ways. For instance, studied content might be erased in an all-or-nothing manner. Alternatively, information may be retained but the precision might be degraded or blurred. Here, we asked whether the similarity of interfering information to memory contents might differentially impact these 2 aspects of forgetting. Observers studied colored images of real-world objects, each followed by a stream of interfering objects. Across 4 experiments, we manipulated the similarity between the studied object and the interfering objects in circular color space. After interference, memory for object color was tested continuously on a color wheel, which in combination with mixture modeling, allowed for estimation of how erasing and blurring differentially contribute to forgetting. In contrast to the similarity assumption, we show that highly *dissimilar* interfering items caused the greatest increase in random guess responses, suggesting a greater frequency of memory erasure (Experiments 1–3). Moreover, we found that observers were generally able to resist interference from highly similar items, perhaps through surround suppression (Experiments 1 and 4). Finally, we report that interference from items of intermediate similarity tended to blur or decrease memory precision (Experiments 3 and 4). These results reveal that the nature of visual similarity can differentially alter how information is lost from memory.

Keywords: interference, visual similarity, visual memory, forgetting, color memory

Interference is a term broadly used to describe the disruption of task-relevant processing by competing information. While interference is commonly associated with forgetting in memory (Conrad, 1964; Keppel & Underwood, 1962; Wixted, 2004; Sadeh, Ozubko, Winocur, & Moscovitch, 2016), past studies have shown that recently encountered information can also bias judgments of immediate perception (Barense et al., 2012; Fischer & Whitney, 2014; Huang & Sekuler, 2010; Newsome, Duarte, & Barense, 2012), and attentional selection (Gao et al., 2016; Olivers, 2009; Soto, Hodsoll, Rotshtein, & Humphreys, 2008). For instance, a build-up of proactive interference in memory can impair perceptual matching of complex objects (Barense et al., 2012; Newsome et al., 2012). Additionally, given that attention is automatically captured by stimuli that match the contents of visual memory (Gao et al., 2016; Olivers, 2009; Soto et al., 2008), interference between items maintained in memory can also impact the deployment of selective attention (Kiyonaga & Egner, 2016). Thus, interference is a process that affects a wide range of cognitive faculties from

---

Sol Z. Sun and Celia Fidalgo, Department of Psychology, University of Toronto; Morgan D. Barense and Andy C. H. Lee, Department of Psychology, University of Toronto, and Rotman Research Institute at Baycrest Hospital, Toronto, Ontario, Canada; Jonathan S. Cant, Department of Psychology, University of Toronto; Susanne Ferber, Department of Psychology, University of Toronto, and Rotman Research Institute at Baycrest Hospital.

Sol Z. Sun and Celia Fidalgo are equal contribution first-authors. Morgan D. Barense, Andy C. H. Lee, Jonathan S. Cant, and Susanne Ferber are equal contribution senior authors.

We thank Timothy Brady for helpful discussion, as well as providing images and MATLAB code for image rotation in color space. The authors also thank Lindsay Arathoon and Idil Askar for assistance in data collection. These data were presented at the annual meeting of the Vision Science Society in 2017. This work was supported by grants from the Natural Sciences and Engineering Research Council (NSERC) awarded to Morgan D. Barense, Andy C. H. Lee, Jonathan S. Cant, and Susanne Ferber, a James S. McDonnell Scholar Award to Morgan D. Barense, and NSERC and Ontario Graduate Scholarships (OGS) awarded to Sol Z. Sun and Celia Fidalgo.

Correspondence concerning this article should be addressed to Sol Z. Sun, Department of Psychology, University of Toronto, 100 South George Street, Toronto, Ontario, Canada, M5S 3G3. E-mail: sol.sun@mail.utoronto.ca

immediate perception to memory over long timescales. However, there is surprisingly little consensus between fields regarding which types of information generate the most interference. In episodic memory research, a widely held assumption is that information highly similar to memory contents generates the most interference, and thus the greatest impairment in performance, which we will refer to as the *similarity assumption* of interference. This intuitive idea is based on earlier studies using verbal stimuli (Keppel & Underwood, 1962; Wickens, 1970), and has been echoed in more recent findings with complex objects (Barense et al., 2012; Makovski, 2016; Watson & Lee, 2013; Yeung, Ryan, Cowell, & Barense, 2013). Conversely, several studies have demonstrated that a greater degree of similarity between studied items can actually benefit memory. In visual change-detection tasks, performance is higher when nontarget studied items are more similar to the target, a finding that holds for low-level visual features such as color, orientation, and line length (Johnson, Spencer, Luck, & Schöner, 2009; Lin & Luck, 2009; Sims, Jacobs, & Knill, 2012), and generalizes to a more complex stimulus class, namely faces (Jiang, Lee, Asaad, & Remington, 2016). These results indicate that greater similarity between studied and interfering information does not necessarily lead to greater performance impairments, as predicted by the similarity assumption. Indeed, recent studies suggest that when nontarget items fall within a critical window of high similarity relative to the contents of memory, processing of these items is attenuated through surround suppression (Kiyonaga & Egner, 2016; Störmer & Alvarez, 2014). Further challenges can be seen in the psychophysics literature on "memory masking." Early studies demonstrated that discrimination thresholds for a single spatial frequency remained constant across extended delays (from a few seconds up to 50 hr, Magnussen & Dyrnes, 1994), but the presentation of a single interfering item during the retention period caused an elevation in threshold, indexing a decrease in performance (Magnussen, Greenlee, Asplund, & Dyrnes, 1991; Nemes, Parry, Whitaker, & McKeefry, 2012; Rademaker, Bloem, De Weerd, & Sack, 2015). In contrast to the similarity assumption, it is often found that performance decreases as the interfering item becomes more *dissimilar* to the studied item (Bennett & Cortese, 1996; LaLonde & Chaudhuri, 2002; McKeefry, Burton, & Vakrou, 2007; Nemes, Whitaker, Heron, & McKeefry, 2011; Rademaker et al., 2015). Taken together, understanding how interference drives the loss of information from memory is clearly important, yet it is difficult to distil previous findings into a set of principles relating target-distractor similarity to interference.

In the experiments reported here, we take steps toward resolving this issue by first assuming that forgetting is not necessarily all-or-nothing, but can be graded. It is possible that after interference, studied information might be erased completely. Alternatively, interference may degrade the precision of studied information, akin to the blurring of representations that are still retained. Furthermore, while a memory manipulation can simultaneously influence both parameters (e.g., directed forgetting, Fawcett, Lawrence, & Taylor, 2016), other manipulations can influence either erasure or blurring, while leaving the other unimpaired (e.g., Dowd, Kiyonaga, Beck, & Egner, 2015; Rajsic, Sun, Huxtable, Pratt, & Ferber, 2016; Sutterer & Awh, 2016;). Thus, we assume that erasure and blurring are, at least to some degree, dissociable aspects of forgetting (Richter, Cooper, Bays, & Simons, 2016). However, because most previous interference studies have tested memory using discrete, forced-choice responses, it is difficult to differentiate between these two alternatives. Indeed, depending on the target-distractor similarity on discrete response tasks, both types of forgetting might manifest as an "incorrect" response. For instance, if interference from highly similar information decreases the precision of memory (i.e., blurring), then such interference would diminish the ability to discriminate the studied item from similar lures. However, highly similar interference could simultaneously act as a stream of imprecise reminders, thereby increasing the likelihood of remembering the *gist* of the studied content. Thus, on the one hand, decreases in memory precision reflect the loss of fine-grained information, despite the potential strengthening of course-grained information. On the other hand, erasure of memory contents would reflect the complete loss of both fine- and course-grained information. However, most forced-choice memory tasks used to date do not vary target-to-lure similarity and thus cannot discriminate between these possibilities.

To overcome this methodological and theoretical challenge, we tested memory using continuous free recall (Wilken & Ma, 2004) which, in conjunction with mixture modeling, allows for the estimation of the separate contributions of erasing and blurring owing to forgetting (Suchow, Brady, Fougnie, & Alvarez, 2013; Zhang & Luck, 2008). Thus, in the present series of experiments, we asked whether the similarity of interfering information relative to memory contents might differentially impact these two aspects of forgetting. Furthermore, by presenting multiple interfering items during the retention period, we also tested whether the similarity *between* interfering items (i.e., the variability of a set of interfering items) might also play a role in interference. It is well documented that observers are sensitive to the variability of a set of items (Brady & Alvarez, 2015; Cant & Xu, 2015; Michael, de Gardelle, & Summerfield, 2014; Tong, Ji, Chen, & Fu, 2015) and that such variability can influence recognition memory decisions (Jiang et al., 2016; Kahana & Sekuler, 2002; Viswanathan, Perl, Visscher, Kahana, & Sekuler, 2010; Yotsumoto, Kahana, Wilson, & Sekuler, 2007). Thus, we investigated how interference might be generated by both the mean similarity of interfering items to the studied item, and the variability between interfering items.

Across all four experiments, participants studied images of real-world colored objects (Figure 1 and Brady, Konkle, Gill, Oliva, & Alvarez, 2013). Recent studies demonstrate that short-term storage capacity is greater for real-world objects than simple stimuli like colored squares (Brady, Störmer, & Alvarez, 2016). Thus we utilized real-world objects as a more naturalistic means of studying visual memory. Furthermore, this allowed us to ask whether observers obligatorily encode conjunctive representations of object identity and color (Experiment 2). Following initial encoding, the color of the object was to be maintained in memory across a retention period where we presented a stream of five interfering colored objects. Between conditions, we manipulated the similarity of the interfering items to the study item, along with the variability between interfering items. Finally, after the retention period, participants recalled the studied color by selecting it on a continuous color wheel.

To preview the results, we found that interference from highly *dissimilar* items resulted in a greater frequency of memory erasure, indexed by an increase in random guess responses (Experiments 1–3). This effect was highly robust, and is not consistent with the

Figure 1. All 75 unique real-world objects used in Experiments 1–4. Object colors (arranged in columns) represent 15 equally spaced steps spanning the entire 360° circular hue space. See the online article for the color version of this figure.

widely held *similarity assumption* that highly similar interfering information necessarily results in the greatest performance impairment. Several control analyses demonstrated that the increase in random guess responses following dissimilar interference cannot be explained by misreports of the interfering items (i.e., "swap" responses; see Between-Experiments Analyses). Additionally, we found that observers are generally able to resist interference from highly similar items, perhaps through surround suppression in feature space (Kiyonaga & Egner, 2016; Störmer & Alvarez, 2014; Experiments 1 and 4). Finally, we found that interference from items of intermediate similarity was most likely to blur memories, thereby decreasing the *precision* or fine-grained aspect of representation (Experiments 3 and 4; for a summary of all experiments, see Figure 13).

General Method

All participants were recruited from the University of Toronto community, reported having normal or corrected-to-normal visual acuity and normal color vision, provided written informed consent in accordance with the APA ethical standards, and received either partial course credit or were paid $12 (CAD) per hour as compensation. All procedures were approved by the University of Toronto Research Ethics Board. For each experiment, we aimed for a sample size of about 25 participants, based on an a priori power analysis conducted using G*Power (Faul, Erdfelder, Buchner, & Lang, 2009). Specifically, assuming a small effect size ($\eta^2 = 0.10$) and a modest mean correlation between repeated measures ($r = .40$), we required 25 participants to achieve a 0.95 probability of detecting a true main effect in a four-condition, one-way repeated measures analysis of variance (ANOVA). These assumptions were validated in the collected data, as the actual effect size and mean correlation between repeated measures exceeded these values. A power level of 0.95 was chosen as this is the point at which p(Type I error) and p(Type II error) are equal at 0.05.

Experiment 1

The first experiment was designed as an initial test of whether similar and dissimilar interfering items would differentially erase or blur the contents of memory. To this end, the colors of the interfering items were either very similar (Proximal) or dissimilar (Distal) to the color held in memory (see Figure 2). The similarity assumption would predict that Proximal interference should cause the greatest degree of impairment, regardless of whether that manifests as more frequent memory erasure, a decrease in precision, or both. Furthermore, the similarity assumption would predict little-to-no effect of Distal interference for either component of forgetting. As an intermediate, we also included a Random interference condition, where interfering colors were sampled broadly throughout color space (see Design and Procedure section). Finally, we contrasted performance in these conditions to a Baseline condition where interfering items were presented in gray scale.

Method

Participants. Twenty-seven individuals were recruited to participate in Experiment 1. Three participants were excluded as a result of poor performance (for a detailed discussion of exclusion criteria, see Exclusion Criteria section). The remaining 24 participants (16 female, 23 right-handed) had a mean age of 18.75 years (range: 18–23 years).

Stimuli and apparatus. The experiment was run using MATLAB (Mathworks, Natick, MA) with the Psychophysics Toolbox (Brainard, 1997; Kleiner et al., 2007; Pelli, 1997) on a 50 × 30 cm LCD monitor with a resolution of 1,920 × 1,080 pixels and a refresh rate of 60 Hz. Viewing distance was held constant at 50 cm and head position was secured with a chinrest.

All stimuli were viewed on a uniform white background. Seventy-five unique images of real-world objects were selected from a database used previously in (Brady, Konkle, Alvarez, & Oliva, 2008, 2013; Figure 1), with each object resized to fit in an invisible frame subtending approximately 7.8° × 7.8° of visual angle. Images were selected such that each object could occur naturally in any color (e.g., a felt pen) and did not have a strong, real-world association with a particular color. On any given trial, study objects were presented in a random color, controlled by a randomly selected angle (bounded by 0°–359°) used to rotate the object in continuous, circular color space (Figure 1; also see Brady

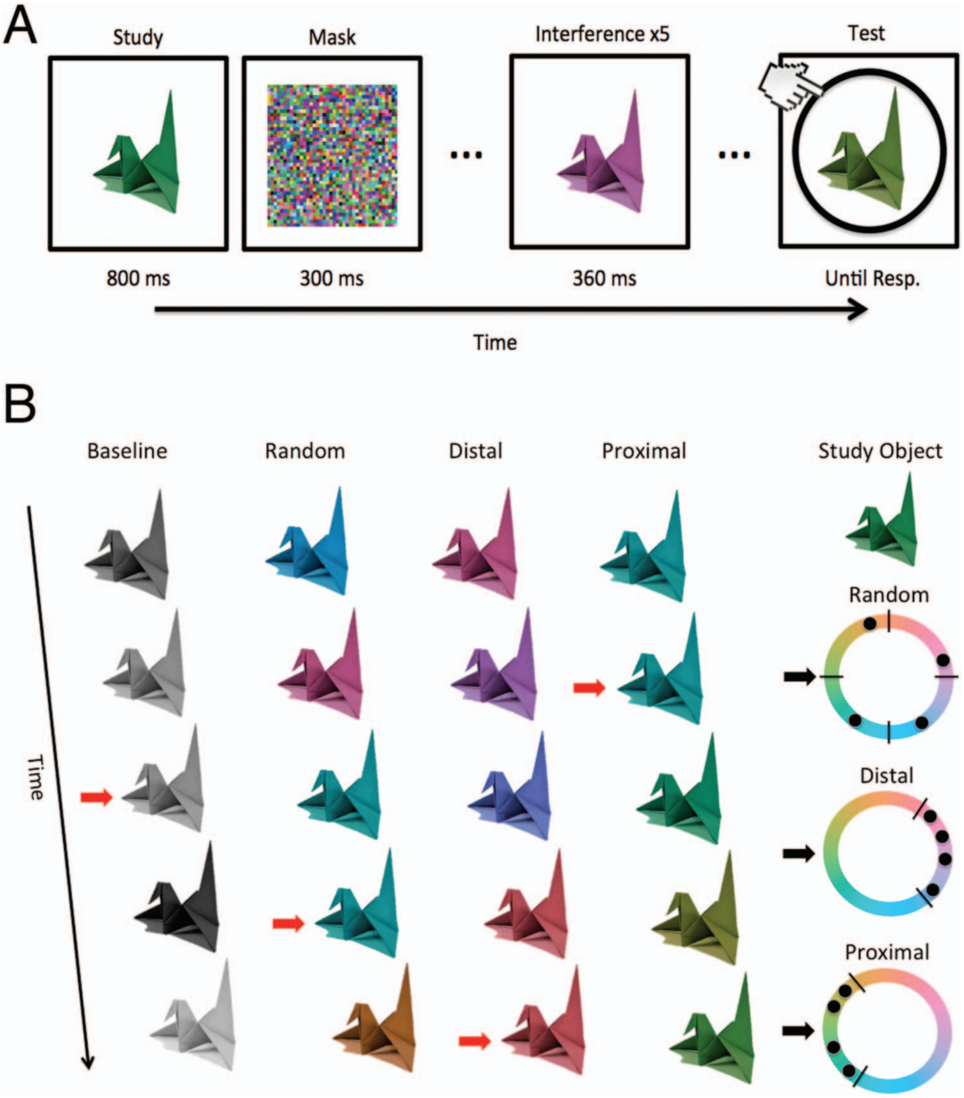

Figure 2. Trial schematic and interference colors in Experiment 1. (A) Example trial schematic for Experiment 1 (initial fixation screen not depicted). (B) Examples of the interference colors. Given the color of the study object (top right) denoted by the black arrow on the color wheels, locations of the interfering colors are denoted as filled black circles, for each condition. Red arrows indicate the repeated color (or luminance) for the one-back task.

et al., 2013). Presentation of study objects was followed by presentation of a mask (subtending ~11.4° × 11.4° visual angle) to prevent visual afterimages. The mask image was generated in MATLAB by filling a 50 × 50 element grid with uniform noise for each RGB value. Interfering objects in Experiment 1 were always of the same identity (e.g., a crane, a tissue box, etc.) as the studied object, but varied in color (or luminance) depending on the condition (see Design and Procedure section).

Design and procedure. Experiment 1 employed a fully within-subject, one-way, four-level design, with a factor of Interference Type (Baseline vs. Random vs. Distal vs. Proximal). Participants completed 75 trials per condition, for a total of 300 trials. Manipulation of Interference Type was mixed (i.e., interleaved), with the sequence of trial types randomized for each participant. Each unique object was presented once per condition, for a total of four repeated instances of the same object identity serving as the studied item in an experimental session. In Experiment 1, interfering objects were always of the same identity as the study object, but varied in their color depending on the condition (see Figure 2).

Participants initiated each trial by pressing the spacebar. Each trial began with the presentation of a central black fixation cross and the word “Remember” for 1,000 ms. The word was presented approximately 1.7° visual angle above fixation in 24 point Geneva font. Next, both the fixation cross and the word were removed and the study object was presented in a random color for 800 ms. Participants were instructed to remember the color of the study item for the duration of the trial. The study object was immediately masked for 300 ms, followed by the sequential presentation of five interfering items, each presented for 360 ms, with a 250 ms interstimulus interval (ISI) between items. To ensure that the interfering items were attended to, the color (or luminance in the Baseline condition) repeated once per stream and participants were instructed to perform a one-back task, pressing the spacebar when the color (or luminance) repeated across consecutive objects. Notably, this task design follows from previous work (Watson & Lee, 2013) demonstrating that a stream of interfering items presented prior to test in the context of a one-back task can be effective at disrupting memory. The repeating color (or luminance) was randomly selected for each trial, as was its location in the sequence. To reduce the chance that participants might stop attending to the interfering items after detecting the repeated item, we additionally instructed participants that there might be two repeated items on an infrequent number of trials (although this was never the case) and to stay attentive throughout the entire interference phase. The interference phase was followed by a 600-ms blank screen, after which participants were tested for their memory of the studied object's color. During the test phase, a probe object of the same identity as the study object was initially presented in gray scale, surrounded by a black circular wheel (radius of 6.8° visual angle). The central probe object was rendered in color upon initial movement of the mouse. Participants were instructed to recall the color of the study object, indicating their choice by moving the mouse around the wheel, thereby smoothly rotating the color of the probe object in circular color space, and clicking the left mouse button to finalize their response. Participants were instructed to prioritize accuracy over speed, thus we focus solely on the former measure in our data analysis. The test phase continued indefinitely until a response was detected (Figure 2A).

As mentioned previously, we manipulated the color similarity between the study object and the interfering objects across conditions. In the Random condition, the entire 360° color wheel was divided into four equal quadrants, each 90° wide, and each of the four unique interfering colors were selected randomly from within each quadrant. It should be noted that while the Random condition presents interference of intermediate similarity, it also incorporates the greatest variability between interfering items. In the Distal condition, the interfering colors were selected from a 120° slice of the color wheel, centered on a point 180° opposite to the angle of the studied item, in color space. This 120° slice was divided into four equal quadrants, each 30° wide, and each of the four interfering colors was drawn randomly from within each quadrant. In the Proximal condition, interfering colors were selected from a 120° slice of the color wheel centered on the angle of the study item. This 120° slice was divided into four equal quadrants, again 30° wide. Each of the four unique interfering colors was chosen randomly from within each quadrant (Figure 2B). For these three conditions, selection of the interfering colors was further constrained such that all colors were separated by a minimum distance of 5° from each other, and the studied color. In the Baseline condition, interfering objects were presented in gray scale, but with varying luminance (range of 0.18–0.82 mean pixel intensity). This condition was intended to emulate the appearance and task demands of the interference conditions, without presentation of values of the to-be-remembered feature (i.e., color) during the interference phase (Figure 2B).

Exclusion criteria. We used hit rates on the one-back task as an indirect measure of whether the interfering items were attended to or not. The one-back hit rate was calculated as the proportion of trials where participants made a response within 1,200 ms of the presentation of the repeated item. Three participants with exceptionally low one-back hit rates were excluded from further analysis in Experiment 1 (hit rates of 1%, 11%, and 57%, respectively). Our exclusion criterion was determined through a simulation, which revealed that we require roughly 40 trials per condition to accurately estimate mixture model parameters (see Appendix A for simulation details; also see Bays, 2016). The remaining 24 participants correctly detected the repeated item on 80.6% of trials, resulting in an average of ~60 trials retained per condition (see Appendix B for a table of descriptive statistics for the one-back task). We conducted all analyses on trials where participants correctly identified the repeating item, thus effectively holding one-back hit rate constant at 100% across all conditions. The results did not qualitatively differ if we did not exclude participants or trials. This was also validated quantitatively, as the test statistics of all analyses conducted on the data without exclusions fell within the 95% bootstrap confidence intervals of the analyses conducted on the data with exclusions. Thus the data without exclusions will not be discussed further.

Analysis. For each trial, we calculated the signed response error as the angular distance between the studied color and the participant's reported color. We fit the distributions of response errors with the two-component mixture model of Zhang and Luck (2008) separately for each participant and for each condition, which allowed us to estimate how erasing and blurring differentially contribute to forgetting. All mixture modeling analyses were performed using MemToolbox (Suchow, Brady, Fougnie, & Alvarez, 2013) in MATLAB. Specifically, the model is given as:

$$p(\hat{\theta}) = (1 - \gamma)\phi_{\sigma}(\hat{\theta} - \theta) + \gamma\frac{1}{2\pi}$$

Variability in the response errors, calculated as the angular distance between the studied color, θ, and the reported color, $\hat{\theta}$, is assumed to be driven by two sources. With a probability of γ, no information about the studied item is retained, and participants guess randomly around the color wheel with responses modeled as a uniform distribution, $\frac{1}{2\pi}$. With a probability of $1 - \gamma$, information about the studied item is retained, but perturbed with noise. These responses are modeled as a Von Mises distribution (ϕ, the circular analogue of the Gaussian distribution for linear data), centered on a mean of 0 (θ, the angle of the studied item), with memory noise represented as circular standard deviation, σ (Implemented in MemToolbox as *StandardMixtureModel*; Figure 3). As mentioned previously, the model is fit to the data separately for each participant and for each condition. The two free parameters, γ (representing guess rate) and σ (*SD*, representing memory noise) are obtained through maximum likelihood estimation (MemToolbox's *MLE* function). To facilitate interpretability, we present all results in terms of precision ($1/\sigma$) and p(Mem), defined as the probability that the studied information is maintained in memory, however precisely ($p(\text{Mem}) = 1 - \gamma$). Higher values on both measures are indicative of better performance.

To assess the reliability of the differences in parameter estimates (p(Mem) and precision) between conditions, for all experiments, we chose to bootstrap all inferential statistical tests to account for violations of statistical assumptions. Indeed, a Kolmogorov–

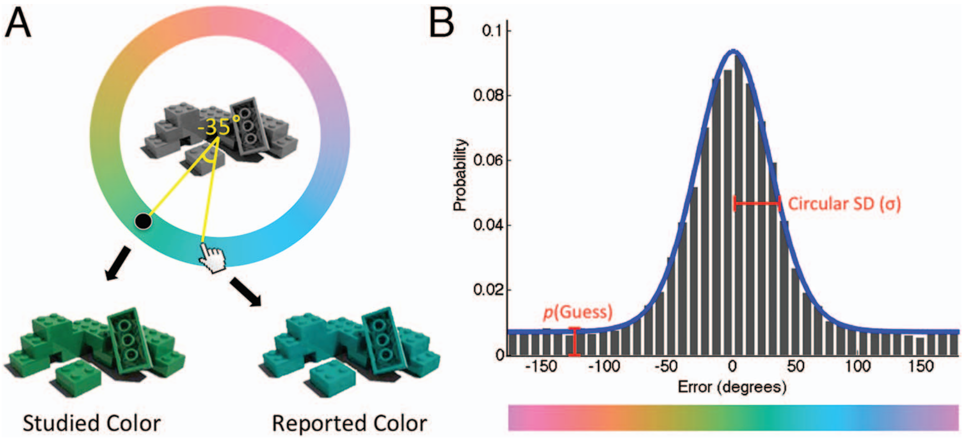

Figure 3. Illustration of continuous free recall and mixture modeling. (A) Signed response error calculated as the angular distance between the studied color and the participant's reported color on a hypothetical trial. (B) Error histogram fit with the two-component mixture model of Zhang and Luck (2008). See the online article for the color version of this figure.

Smirnov (KS) test revealed that the distribution of p(Mem) deviated from that of a normal distribution, and a Mauchly's test revealed that the between-condition correlations of precision violated sphericity assumption of the repeated measures ANOVA. Thus for any given statistical test (i.e., t, F), we ran 5,000 bootstrap resamples. For each iteration, we resampled the data with replacement, and then calculated both the test statistic and its effect size. For all ANOVA main effects and interactions, we then calculated a discrete p value from the resulting F distribution by taking the proportion of bootstrap resamples that fell below the null value of 1 as an approximation to the more commonly reported parametric p value. After each ANOVA, we compared each interference condition (e.g., Random, Distal, and Proximal) to the Baseline condition using follow-up t tests. If a condition was significantly different from Baseline, we further compared this condition with the other interference conditions. For all t tests, discrete p values were calculated by taking the proportion of bootstrap resamples falling below a null value of 0, then multiplying this proportion by 2, as all comparisons were two-tailed. It should be noted that unlike their parametric counterparts, discrete p values could take on values of exactly 0 or 1 as they are not based on an assumed continuous distribution, but rather a discrete frequency distribution of bootstrap resamples. For effect sizes (i.e., η^2, Cohen's d), we calculated 95 percentile confidence intervals (CI) to assess their reliability (Efron, 1992). Unlike parametric CIs, percentile bootstrap CIs could be asymmetrical in shape, as they are calculated from empirical estimates of the distribution of the statistic, rather than an assumed distribution (Preacher & Hayes, 2008). Our calculation of percentile CIs and discrete p values is identical to the implementation in the "mediation" package for R (Tingley, Yamamoto, Hirose, Keele, & Imai, 2014). All inferential analyses were performed in R (R Core Team, 2014) using custom scripts and the "car" package (Fox & Weisberg, 2011).

Results and Discussion

Differences in p(Mem) and precision between interference conditions were tested using one-way repeated measures ANOVAs with a factor of Interference Type (Baseline vs. Random vs. Distal vs. Proximal). The analysis revealed that p(Mem) reliably differed between conditions, $F(3, 69) = 8.53$, $p = .001$; $\eta^2 = 0.27$, 95% CI [0.13, 0.46] (Figure 4A–B). To determine which types of interference were most likely to erase information from memory, we compared p(Mem) in each of the interference conditions to the Baseline condition using paired t tests. To this end, we observed a reliable decrease in p(Mem) in the Distal condition, relative to Baseline, $t(23) = 4.45$, $p < .01$; $d = 0.91$, 95% CI [0.59, 1.36]. Furthermore, Distal interference also decreased p(Mem) relative to both Random, $t(23) = 2.13$, $p = .02$; $d = 0.43$, 95% CI [0.07, 0.85] and Proximal, $t(23) = 3.61$, $p < .01$ $d = 0.74$, 95% CI [0.40, 1.14] interference. Thus, in contrast to the similarity assumption of interference, these results suggest that distance in feature space does not shield information from interference, but on the contrary, increases the probability that the originally studied information cannot be reported. Furthermore, while we did not observe a reliable decrease in p(Mem) in the Random condition, relative to Baseline, $t(23) = 1.73$, $p = .10$; $d = 0.35$, 95% CI [−0.03, 0.66], we do observe a reasonable numerical difference, perhaps due to the fact that two of four unique interfering colors in the Random condition could be quite dissimilar to the studied color (also see Results section, Experiment 2). Finally, p(Mem) did not show a decrease in the Proximal condition, $t(23) = 1.16$, $p = .23$; $d = 0.24$, 95% CI [−0.17, 0.62], relative to Baseline. These results suggest that the main effect from the omnibus ANOVA was driven primarily by a substantial decrease in p(Mem; i.e., an increase in random guess responses) following interference by information highly dissimilar to the contents of memory.

Turning now to precision, calculated as $1/\sigma$ estimated from the fitted model, we also found reliable differences between interference conditions, $F(3, 69) = 3.63$, $p = .016$; $\eta^2 = 0.14$, 95% CI [0.05, 0.30] (Figure 4C–D). More specifically, memory representations in the Random condition were considerably less precise, relative to Baseline, $t(23) = 2.32$, $p = .014$; $d = 0.47$, 95% CI [0.09, 0.93]. Furthermore, Random interference also decreased precision relative to both Distal, $t(23) = 2.84$, $p = .004$; $d = 0.58$, 95% CI [0.16, 1.27],

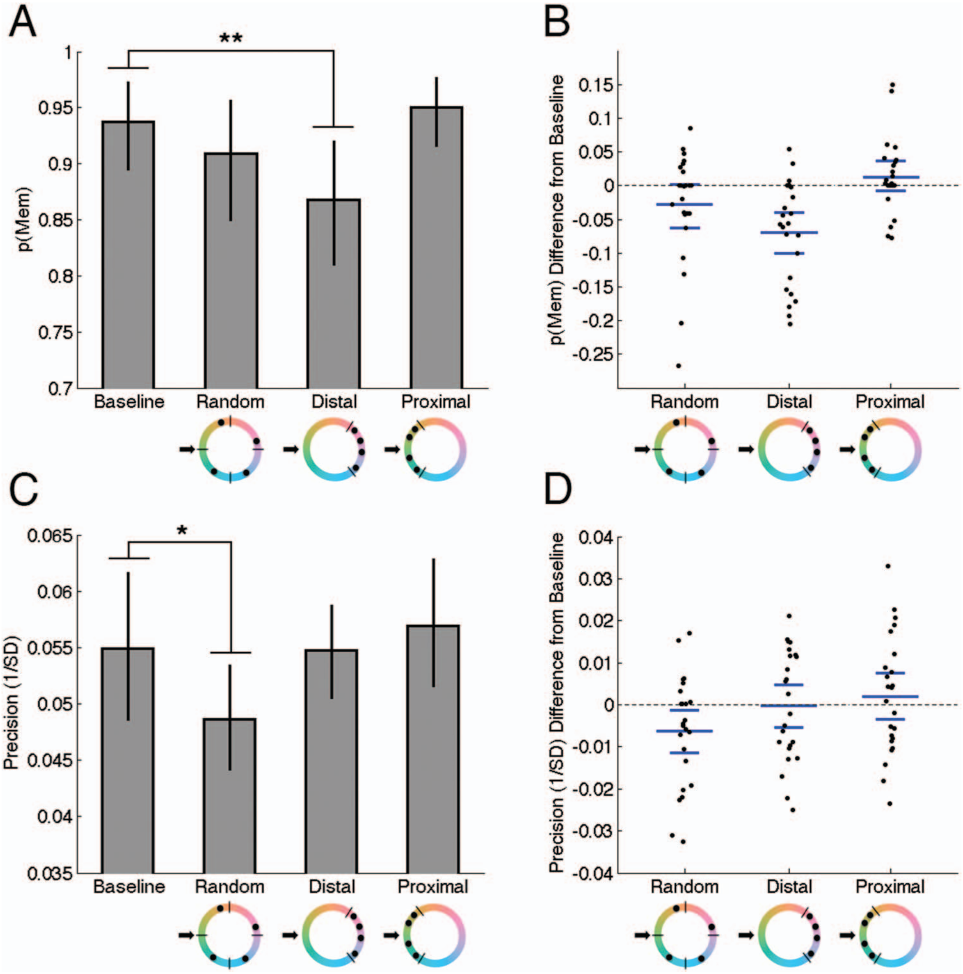

Figure 4. Results of Experiment 1. (A) p(Mem), the probability that the studied information is retained in memory, however precisely, across the interference conditions. Error bars represent 95% CIs calculated from bootstrap distributions assessing the reliability of each condition mean. (B) The interference effect in p(Mem), calculated as the difference in p(Mem) between an interference condition and the Baseline condition for individual participants. The hashed line at 0 represents the mean of the Baseline condition. The central blue line denotes the condition mean, while the outer blue lines denote the 95% bootstrap confidence interval. Any given condition is significantly different from Baseline if the hashed line at 0 falls outside of the 95% CI denoted by the outer blue lines. (C–D) Differences in memory precision between conditions, and the interference effect plotted in the same manner as (A–B), respectively. (* $p < .05$. ** $p < .01$). See the online article for the color version of this figure.

and Proximal, $t(23) = 2.93$, $p < .01$; $d = 0.60$, 95% CI [0.27, 1.08] interference. Interestingly, we did not observe a decrease in precision in either the Distal, $t(23) = 0.05$, $p = .96$; $d = 0.01$, 95% CI [−0.43, 0.44], or the Proximal conditions, relative to Baseline, $t(23) = 0.70$, $p = .52$; $d = 0.14$, 95% CI [−0.30, 0.55]. Tentatively, these results seem to suggest that either high variability between interfering items, or interference of *intermediate* similarity causes decreases in the fidelity or precision of memory representations (see Experiments 3 and 4). In sum, counter to the predictions of the similarity assumption that highly similar interference should result in the greatest performance impairment, our results highlight that the nature of visual similarity differentially alters the way in which information is lost from memory.

Experiment 2

The results of Experiment 1 showed that dissimilar interference increased the probability of random guess responses (i.e., decreased p(Mem)), thus indexing an increased likelihood that the originally studied information was erased from memory. Conversely, highly variable interference of intermediate similarity decreased memory precision. Surprisingly, we did not find any effect of similar interfering information (but see Experiments 3 and 4). As our results (particularly in the Distal condition) are not consistent with the widely held similarity assumption, one goal of Experiment 2 was to replicate and assess the reliability of the effects observed in Experiment 1. Furthermore, we also looked to test whether observers would obligatorily encode a conjunctive representation of both object color and identity even though object identity was task-irrelevant. One recent study found that both an item's shape and color are automatically encoded, and both features can subsequently guide attention (Gao et al., 2016). Other studies have reported that while multiple item features are obligatorily encoded into visual memory, their maintenance is voluntary, and one feature may be dropped prior to test (Marshall & Bays, 2013). Thus, in Experiment 2, we manipulated the identity of the interfering items, as well as their color relative to the studied item. This was motivated by the *representational hierarchical model* (Cowell, Bussey, & Saksida, 2010; Graham, Barense, & Lee, 2010), which posits that object representations will be more resistant to interference if multiple features (e.g., color and shape) are stored as a conjunction, which reduces feature-level overlap with interfering objects (Bartko, Cowell, Winters, Bussey, & Saksida, 2010). Thus, we reasoned that if observers encoded the real-world object's identity (e.g., a t-shirt, a felt pen) and its color as a conjunction, then interference would be reduced if interfering colors were presented on a different object. However, if observers maintained only a representation of color, then the identity of the interfering objects should not interact with the color interference effects found in Experiment 1.

Method

Participants. Twenty-eight individuals were recruited to participate in Experiment 2. Three participants were excluded for poor performance on the one-back task (hit rates of 37%, 53%, and 57%, respectively), based on the same exclusion criterion as used in Experiment 1 (see Experiment 1 Exclusion Criteria section). The remaining 25 participants (14 female, 23 right-handed) had a mean age of 20.52 years (range: 18–28 years).

Design and procedure. Experiment 2 employed a fully within-subject, 3 × 2 design, with factors of Interference Type (Baseline vs. Random vs. Distal) and Distractor Identity (Same vs. Different). Given that we added an additional manipulation, for this experiment we removed the Proximal condition to ensure that we could run a sufficient number of trials per condition in a single experimental session (but see Experiment 3). Participants completed 60 trials of each condition, for a total of 360 trials. All trial types were presented in a mixed (i.e., interleaved), random sequence for each participant. For all trials, the five interfering items were of the same identity. However, on half of all trials, the identity of the interfering objects was the same as the studied object, as in Experiment 1. For the remaining trials, the identity of the interfering objects was different from the studied object, randomly selected with the constraint that no object identity would serve as the interfering object more than once per Interference Type (see Figure 5). All other procedures were identical to Experiment 1.

Results and Discussion

Differences in p(Mem) and precision between conditions were tested using 3 × 2 repeated measures ANOVAs with factors of Interference Type (Baseline vs. Random vs. Distal) and Distractor Identity (Same vs. Different). All inferential tests were evaluated

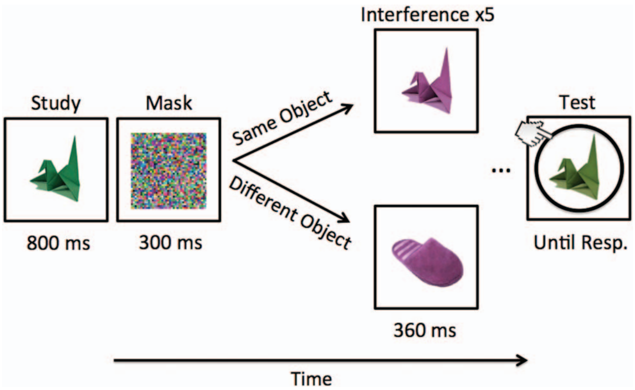

Figure 5. Example trial schematic for Experiment 2. For half of all trials, the identity of the five interfering objects differed from the studied object. See the online article for the color version of this figure.

using the same bootstrap procedure as Experiment 1. We found that p(Mem), the probability of maintaining the studied information in memory, no matter how precisely, differed depending on the color of the interfering items, leading to a reliable main effect of Interference Type, $F(2, 48) = 4.79$, $p = .019$; $\eta_p^2 = 0.17$, 95% CI [0.05, 0.37] (Figure 6A–B). However, p(Mem) did not differ depending on the identity of the interfering objects, and thus we did not observe a reliable main effect of Distractor Identity, $F(1, 24) = 0.10$, $p = .75$; $\eta_p^2 = 0.004$, 95% CI [0, 0.20], or an interaction between the two factors, $F(2, 48) = 0.44$, $p = .65$; $\eta_p^2 = 0.02$, 95% CI [0, 0.18]. These results were mirrored in precision such that a reliable main effect of Interference Type was observed, $F(2, 48) = 5.61$, $p = .005$; $\eta_p^2 = 0.19$, 95% CI [0.07, 0.39], with no main effect of Distractor Identity, $F(1, 24) = 0.03$, $p = .87$; $\eta_p^2 = 0.001$, 95% CI [0, 0.21], and no interaction, $F(2, 48) = 1.40$, $p = .31$; $\eta_p^2 = 0.06$, 95% CI [0.002, 0.29] (Figure 6C–D). Thus, we did not find evidence suggesting that the effects of color interference might differ depending on whether the identity of the interfering objects matched the identity of the studied object (see General Discussion section).

Given that the identity of interfering objects did not affect either memory parameter, we collapsed across this factor and compared the Random, Distal, and Baseline conditions using paired t tests. To this end, we observed the largest decrease in p(Mem) in the Distal condition, relative to Baseline, $t(24) = 2.83$, $p = .002$; $d = 0.57$, 95% CI [0.23, 0.96]; Figure 6B, again suggesting that highly dissimilar items increase the probability that the originally studied information will be erased from memory, replicating the results of Experiment 1. We also found a reliable decrease in p(Mem) in the Random condition, relative to Baseline, $t(24) = 2.22$, $p = .014$; $d = 0.44$, 95% CI [0.11, 0.78], a trend that was also observed in Experiment 1. Furthermore, given that the effect of Random interference on p(Mem) was greater here than in Experiment 1, we did not find a difference between Random and Distal interference in the present experiment, $t(24) = 1.02$, $p = .32$; $d = 0.20$, 95% CI [−0.18, 0.66]. With respect to precision, we again found a reliable decrease in the Random condition, relative to both Baseline, $t(24) = 3.06$, $p = .002$; $d = 0.61$, 95% CI [0.24, 1.16]; Figure 6D and Distal, $t(24) = 2.91$, $p < .001$; $d = 0.58$, 95% CI [0.29, 0.95] interference. As in Experiment 1, precision in the Distal

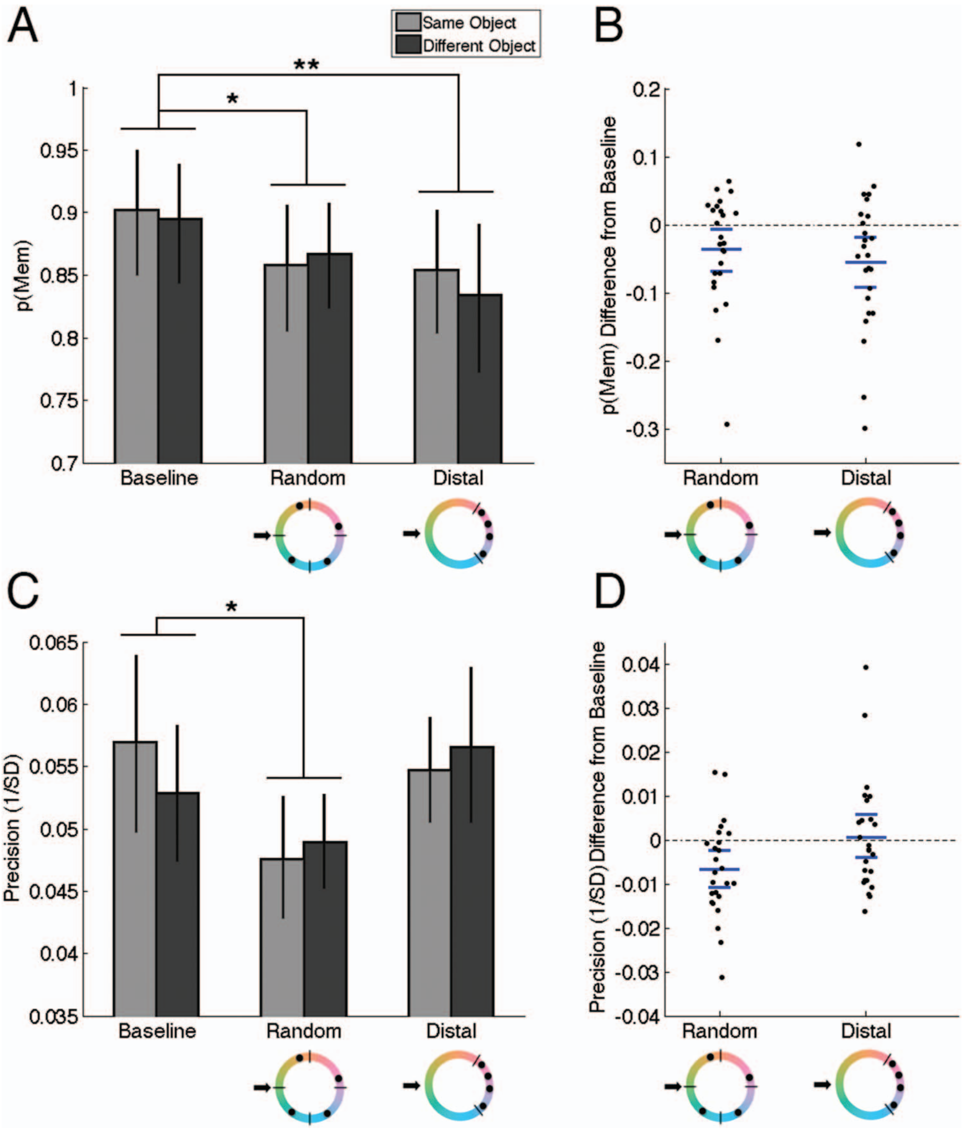

Figure 6. Results of Experiment 2. (A) p(Mem), the probability that the studied information is retained in memory, however precisely, across the interference conditions. Error bars represent 95% CIs calculated from bootstrap distributions assessing the reliability of each condition mean. (B) The interference effect in p(Mem), calculated as the difference in p(Mem) between an interference condition and the Baseline condition for individual participants. The hashed line at 0 represents the mean of the Baseline condition. The central blue line denotes the condition mean, while the outer blue lines denote the 95% bootstrap confidence interval. Any given condition is significantly different from Baseline if the hashed line at 0 falls outside of the 95% CI denoted by the outer blue lines. (C–D) Differences in memory precision between conditions, and the interference effect plotted in the same manner as (A–B), respectively. ($^{*}p < .05$. $^{**}p < .01$). See the online article for the color version of this figure.

condition did not differ from Baseline, $t(24) = 0.29$, $p = .77$; $d = 0.06$, 95% CI [−0.45, 0.40]. Taken together, we replicated two key findings from Experiment 1, namely that Random interference causes a decrease in the precision of memory, while Distal interference items increase the likelihood of memory erasure.

Experiment 3

In Experiment 3, we looked to further examine why Proximal interference did not decrease either p(Mem) or precision, when such an effect is predicted by the results of many previous studies (Barense et al., 2012; Conrad, 1964; Fidalgo, Changoor, Page-Gould, Lee, & Barense, 2016; O'Neil, Watson, Dhillon, Lobaugh, & Lee, 2015; Yeung et al., 2013). To this end, we reasoned that perhaps Proximal interference in Experiment 1 was resisted through surround suppression in feature space. Recent studies have shown that attending to a particular value in feature space (e.g., a specific shade of green) amplifies the processing of that value, while attenuating the processing of similar, but nonidentical values (Störmer & Alvarez, 2014).

Furthermore, this effect extends beyond attention to immediate visual stimuli, and also holds for information maintained in memory (Kiyonaga & Egner, 2016). Surround suppression mechanisms have also been implemented in computational models of visual short-term memory (STM; Johnson et al., 2009), and have received support from fMRI studies, notably in visual area V4, which plays a critical role in color perception (Sneve, Sreenivasan, Alnæs, Endestad, & Magnussen, 2015). Based on these findings, we reasoned that Proximal interfering items in Experiment 1 were perhaps too similar to the studied color, and thus attenuated through surround suppression. Thus, in Experiment 3, we increased the variability of the Proximal and Distal conditions (see Experiment 3 Design and Procedure section). We predicted that increasing the variability of interfering items in the Proximal condition would decrease memory precision. We also predicted that we should again replicate a decrease in p(Mem) following Distal interference, with no effect on p(Mem) following Proximal interference.

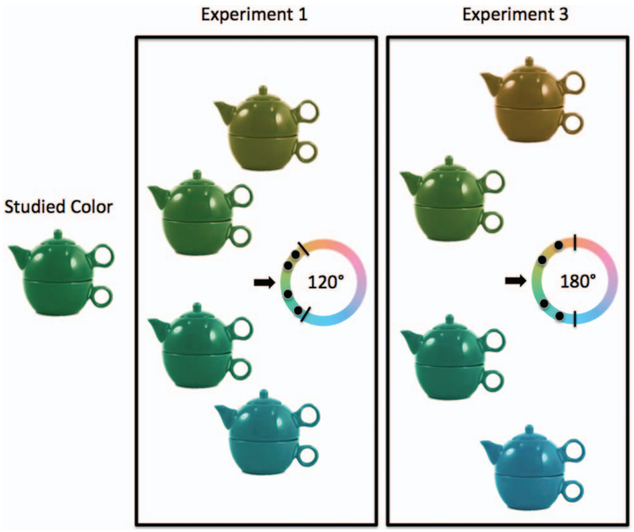

Figure 7. Illustration of the difference in interfering color variability between Experiment 1 and Experiment 3 for the Proximal condition. In Experiment 3, Distal and Proximal interfering colors are drawn from half of the color wheel (180°, previously 120°) either centered opposite to, or on the studied color, respectively. An example studied color for a hypothetical Proximal trial is shown on the left. The central and rightmost panes contain representative examples of interfering items from Experiments 1 and 3, respectively. See the online article for the color version of this figure.

Method

Participants. Twenty-nine individuals (19 female, 24 right-handed) with a mean age of 20.62 years (range: 18–28 years) were recruited to participate in Experiment 3. Based on the same exclusion criterion used in Experiments 1 and 2 (see Experiment 1–Exclusion Criteria), we did not exclude any participants.

Design and procedure. Experiment 3 employed a fully within-subject, one-way, three-level design, with a factor of Interference Type (Baseline vs. Distal vs. Proximal). Participants completed 75 trials of each condition, for a total of 225 trials. All trial types were presented in a random sequence for each participant. Furthermore, as in Experiment 1, on all trials the identity of the interfering object was the same as the studied object. In Experiment 3, we aimed to increase the variability of interfering items in the Distal and Proximal conditions (see Figure 7). To this end, for the Distal condition, we selected a 180° portion of the color wheel (previously 120°) located 180° opposite to the angle of the studied color (i.e., the opposite half of the wheel). We then partitioned this opposite half into four equal slices (each 45° wide) and selected one unique interfering color from within each slice. For the Proximal condition, we repeated this procedure using a 180° portion of the wheel (previously 120°) centered on the studied color and partitioned the half into four equal quadrants. Each interfering color was selected from one of the four quadrants. As in the previous experiments, selection of all interfering colors was constrained such that the interfering colors were separated from each other, and the studied color, by at least 5°. All other procedures were identical to Experiment 1.

Results and Discussion

Differences in p(Mem) and precision between conditions were tested using one-way, repeated measures ANOVAs with a factor of Interference Type (Baseline vs. Distal vs. Proximal). We found that p(Mem) differed between conditions, resulting in a reliable main effect of Interference Type, $F(2, 56) = 10.07$, $p < .001$; $\eta^2 = 0.26$, 95% CI [0.12, 0.45] (Figure 8A–B). Replicating the results of the two previous experiments, we found that highly dissimilar interference increased the probability of the originally studied information being erased from memory, as p(Mem) was lower in the Distal condition, relative to both Baseline, $t(28) = 2.62$, $p = .005$; $d = 0.49$, 95% CI [0.14, 0.88] and Proximal, $t(28) = 4.20$, $p < .01$; $d = 0.78$, 95% CI [0.52, 1.14] conditions. Furthermore, we found that Proximal interference did not reliably impact p(Mem), relative to Baseline, $t(28) = 1.83$, $p = .07$; $d = 0.34$, 95% CI [−0.01, 0.76], with a trend in the opposite direction, which was also observed in Experiment 1.

For precision, the effect of Interference Type at the omnibus level was not reliable, with a small effect size, $F(2, 56) = 2.73$, $p = .08$; $\eta^2 = 0.09$, 95% CI [0.01, 0.29] (Figure 8C–D). Furthermore, we did not observe a reliable decrease in precision in the Distal condition, relative to Baseline, $t(28) = 0.70$, $p = .48$; $d = 0.13$, 95% CI [−0.24, 0.54]. However, increasing the variability between interfering items did lead to a reliable reduction in precision in the Proximal condition, relative to Baseline at the pairwise test level, $t(28) = 2.21$, $p = .03$; $d = 0.41$, 95% CI [0.04, 0.93] (Figure 8D). Furthermore, Proximal interference did not decrease precision relative to Distal, $t(28) = 1.64$, $p = .08$; $d = 0.30$, 95% CI [−0.04, 0.65]. Finally, to directly assess whether increasing the variability of Proximal interfering items decreases precision, we require a direct test between the Proximal conditions of Experiments 1 and 3. Although this between-subjects test affords lower statistical power, we find a marginal difference in precision when the variability of the interfering items is increased (parametric two-sample t test: $t(51) = 2.00$, $p = .05$; $d = 0.55$).

In summary, we found that Distal interference decreases p(Mem) relative to both Baseline and Proximal interference, replicating results of the previous experiments. Furthermore, with an increase to the variability of the interfering items, Proximal inter-

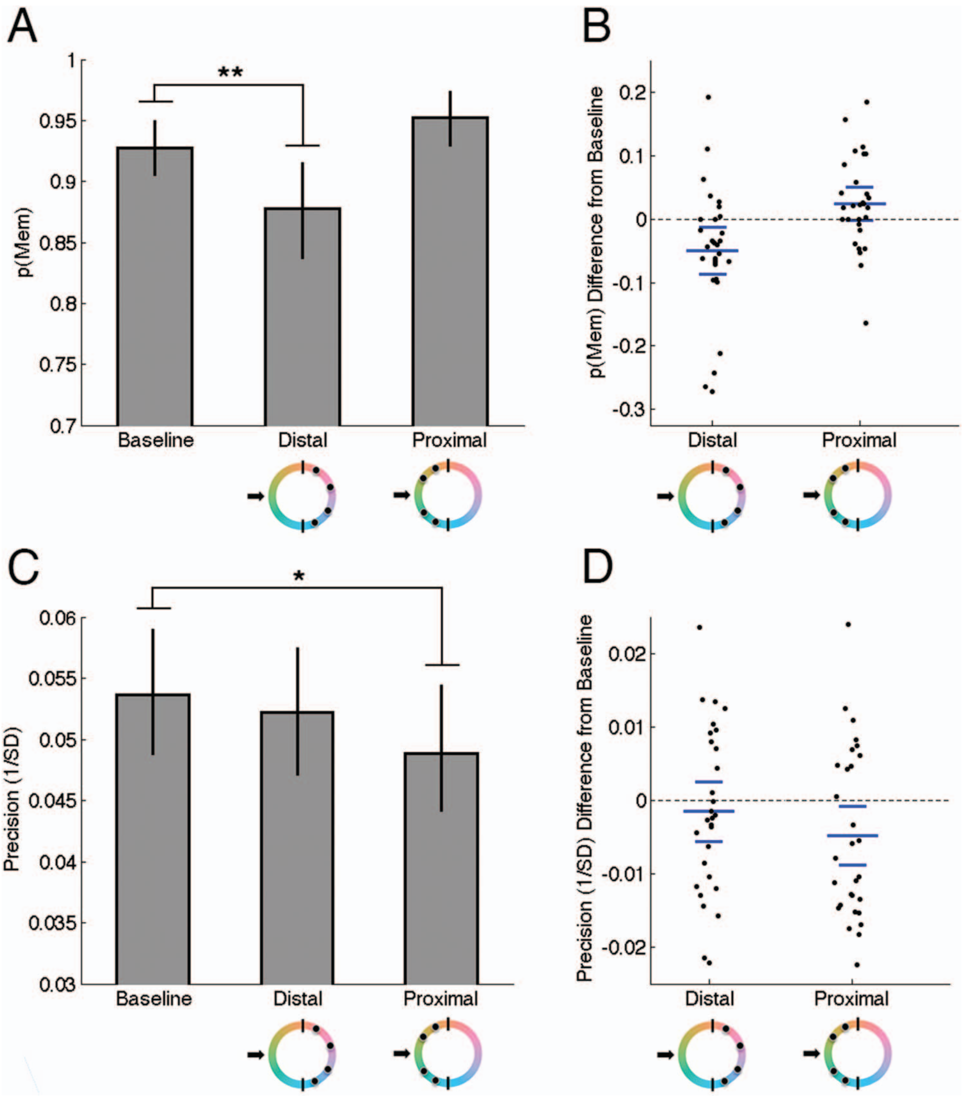

Figure 8. Results of Experiment 3. (A) p(Mem), the probability that the studied information is retained in memory, however precisely, across the interference conditions. Error bars represent 95% CIs calculated from bootstrap distributions assessing the reliability of each condition mean. (B) The interference effect in p(Mem), calculated as the difference in p(Mem) between an interference condition and the Baseline condition for individual participants. The hashed line at 0 represents the mean of the Baseline condition. The central blue line denotes the condition mean, while the outer blue lines denote the 95% bootstrap confidence interval. Any given condition is significantly different from Baseline if the hashed line at 0 falls outside of the 95% CI denoted by the outer blue lines. (C–D) Differences in memory precision between conditions, and the interference effect plotted in the same manner as (A–B), respectively. (* $p < .05$. ** $p < .01$). See the online article for the color version of this figure.

ference decreased precision relative to Baseline, but not relative to Distal interference. Despite the consistency of both effects, the main driving factor behind them remains unclear. With regard to decreases in p(Mem) in the Distal condition, it may be that our results are explained by misreports (i.e., swap errors) of the interfering items. We begin addressing this issue in the following section. With regard to decreases in precision in the Proximal condition, high variability between interfering items could be the key factor that leads to precision declines. As seen in Experiments 1–2, Random interference leads to a robust decrease in precision, either due to high variability between interfering items, or interference of intermediate similarity. The latter is, in fact, consistent with a surround suppression account. Specifically, if maintaining a color in memory causes suppression of similar, but nonidentical colors, then highly similar interference should be attenuated (Kiyonaga & Egner, 2016; Störmer & Alvarez, 2014). However, a precision decrease might emerge as interference falls outside this suppressive window. Thus, in Experiment 4, we looked to test these two accounts by varying the mean dissimilarity of the interfering items from the study item without changing their variability.

Between-Experiments Analyses: Memory Erasure is not Accounted for by Swap Errors

When cued to report a target item, observers will sometimes report a nontarget item instead. These misreports of nontarget items (a.k.a., swap errors, binding errors) are typically assessed using the three-component mixture model developed by Bays and colleagues (Bays, Catalao, & Husain, 2009, implemented in MemToolbox as *SwapModel*). This model is similar to the Zhang and Luck (2008) model used throughout the article thus far, but with an additional component that models reports of nontarget colors (i.e., interfering colors in our task) with the same Von Mises noise as the target distribution. When interpreting decreases in memory precision following Proximal interference (Experiment 3), we reasoned that examining swap errors was not feasible. Specifically, an imprecise target representation may result in the inability to distinguish the target from a perceptually similar interfering item. During the subsequent test, if a response lands close to an interfering color, we cannot elucidate whether this was a true misreport or whether this was the observer's best estimate of an intact, but blurred, target representation. Put differently, when interfering items are highly similar, imprecise target responses might empirically manifest as swap errors, and conversely, swap errors might manifest as decreases in precision. For this reason, swap errors and decreased precision are conceptually intertwined when considering highly similar interference. Additionally, previous work has also shown that swap errors are most likely to occur when nontargets are highly similar to a given target (e.g., Schneegans & Bays, 2017). As a result, the distribution of swap errors and the distribution of target responses are likely highly overlapping, and difficult to distinguish. Nevertheless, we can proceed to examine whether memory erasure following Distal interference is accounted for by swap errors.

In all three experiments thus far, we found that interference from dissimilar information in the Distal condition was most likely to decrease *p*(Mem), suggesting more frequent erasure of the originally studied item. We contend that Distal interference increases the probability of "sudden death," where item representations terminate completely, with a loss of both fine- and course-grained information (Zhang & Luck, 2009). Following sudden death of the studied item, it is assumed that participants guess randomly around the color wheel, with responses distributed uniformly. However, an alternative explanation is that participants may systematically report the color of the interfering items. To examine this possibility, we begin by formally comparing the fit of the *StandardMixtureModel* (Zhang & Luck, 2008) and the *SwapModel* (Bays et al., 2009) to data from the Distal condition for all participants across Experiments 1–3. This analysis was restricted to the Distal condition because we have observed the most reliable decreases in *p*(Mem) in this condition. For simulations justifying whether the three-component *SwapModel* (Bays et al., 2009) is appropriate to use when nontarget colors are not independently sampled, see Appendix C, Figure A2. Figure 9A plots differences in Bayesian information criterion (BIC; Schwarz, 1978), where each dot is a participant, and positive values denote evidence that a given participant is misreporting interfering colors with a non-negligible frequency (i.e., support for the *SwapModel*). This analysis reveals that following Distal interference, the majority of participants do not frequently misreport an interfering color, as the *StandardMixtureModel* was preferred in 76 out of 78 participants.

Next, we looked to select a subsample of participants who committed swap errors at a near-negligible frequency (<1% of trials). To this end, we removed the 18 participants with the strongest *SwapModel* preference and examined how Distal interference affects *p*(Mem) and precision in the remaining 60 partic-

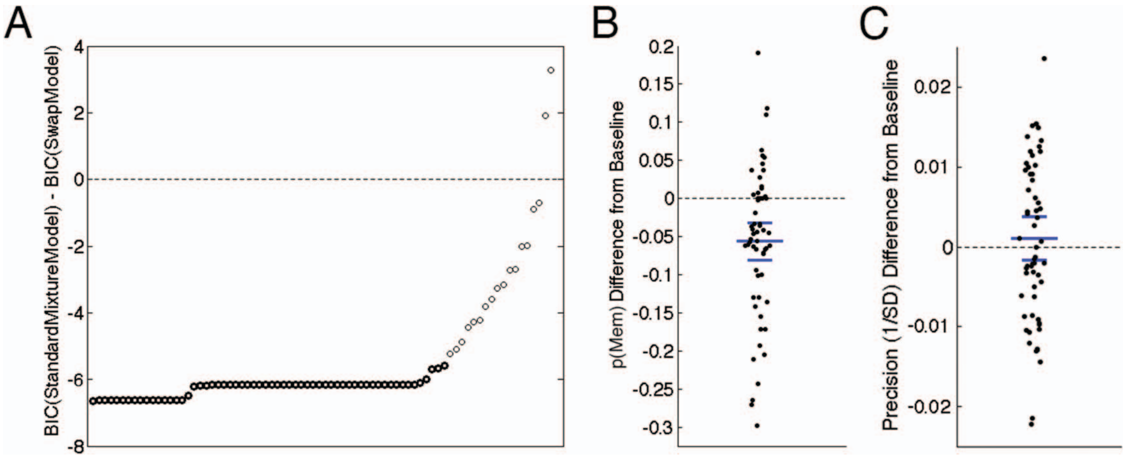

Figure 9. Memory erasure is not accounted for by swap errors. (A) Bayesian information criterion (BIC) differences comparing the 2-component *StandardMixtureModel* with the 3-component *SwapModel* for the Distal condition across Experiments 1–3. Dots represent individual participants. Values represent the likelihood of a model, given the data, and competing models. Positive values indicate support for the *SwapModel*. Bolded dots represent a subsample of participants that commit swap errors on less than 1% of trials. (B and C) Analysis of memory parameters in the subsample of participants with <1% swap errors. Dots represent the difference in *p*(Mem; B), or precision (C) between Distal and Baseline. The central blue line denotes the mean, while the outer blue lines denote the 95% bootstrap confidence interval. The hashed line at zero denotes the mean of the Baseline condition. See the online article for the color version of this figure.

ipants, relative to Baseline. In Figure 9A, bolded dots highlight the 60 participants that were retained after exclusion. For this group, the mean swap rate estimated from the model was less than 1% of trials, and the *StandardMixtureModel* was roughly six times more likely than the *SwapModel*. In this group, we still observe a robust decrease in p(Mem) relative to Baseline, $t(59) = 4.5$, $p < .01$; $d = 0.58$, 95% CI [0.34, 0.85]; Figure 9B, with no change in precision, $t(59) = 0.70$, $p = .49$; $d = 0.10$, 95% CI [−0.17, 0.33]; Figure 9C. For comparison, when all 78 participants were considered, the reduction in p(Mem) following Distal interference was 5.7% (Cohen's $d = 0.62$), relative to Baseline. For the subsample of 60 participants who swap on <1% of trials, the reduction was 5.6% (Cohen's $d = 0.58$). Thus the exclusion of participants who were more likely to commit swap errors only slightly decreased the Distal effect size, suggesting that memory erasure following Distal interference was fairly constant throughout the entire sample, regardless of the frequency of swap errors. Importantly, this demonstrates that Distal interference decreases p(Mem), even when misreports of the interfering colors are accounted for.[1]

Our next analysis focused on data from the Distal condition, exclusively from Experiments 1 and 2, where the interfering colors were selected from a 120° slice of the color wheel. In these data, if participants were systematically reporting the interfering colors following Distal interference, then responses should be asymmetrically clustered around a region of color space centered ~180° opposite to the angle of the studied item, as this region includes the colors of Distal interfering items. Figure 10A plots a histogram of the absolute error of all trials from the Distal condition, across all participants in Experiments 1 and 2. For reference, Figure 10A is the unsigned analogue to the illustrative histogram of signed errors plotted in Figure 3B. In Figure 10A, the x-axis is partitioned into three regions. The leftmost region (faded gray bars, 0°–59°) primarily represents responses that contribute to p(Mem), that is, target responses. The central region of Figure 10A (gray bars, 60°–119°) primarily represents guess responses that do not overlap with the color of Distal interfering items, while the rightmost region (black bars, 120°–179°) primarily represents guess responses that fall in the same region of color space as the Distal interfering items. As mentioned previously, if guess responses in the Distal condition are systematically biased toward the color of the interfering items, then we should observe a greater frequency of responses in the rightmost region, relative to the central region of Figure 10A. Contrary to this prediction, the data appear to be uniformly distributed across the central and rightmost regions, and in fact, we observed slightly more responses in the central region (175 responses), relative to the rightmost region (164 responses). A χ^2 test revealed that the frequency of responses did not significantly differ between the central and rightmost regions ($\chi^2 = 0.71$, $p = .40$).

Finally, we ran a series of simulations to determine the lowest p(Swap) rate at which the frequency of responses would significantly differ between the central and rightmost regions. Specifically, given that the estimated p(Guess) across all Distal trials was 0.14, we varied the simulated p(Swap) value from 0−0.14 in steps of 0.005. For each p(Swap) bin, we calculated a χ^2 comparing the frequency of responses between the central and rightmost bins, as we did in the empirical data. Furthermore, the data were simulated with the same memory precision and number of trials as the empirical data. For each p(Swap) bin, we simulated 10 data sets and report the median χ^2 p value for each bin (Figure 10B). This analysis revealed that the frequency of responses would significantly differ between the central and rightmost bins if observers were misreporting a Distal interfering color (with noise) on ~2% of trials or more. This suggests that our aforementioned analysis has sufficient sensitivity to detect misreports of the interfering items. This analysis further suggests that in the empirical data, misreports of the Distal interfering items occurred on less than 2% of trials. Together with the model comparison analysis, we can also conclude that we would see the same effects of Distal interference on p(Mem) even if the swap error trials are excluded from analysis. Taken together, these analyses suggest that misreports of the interfering items cannot explain decreases in p(Mem) following Distal interference.

Experiment 4

Experiments 1 and 2 showed that Random interference from colors sampled broadly throughout color space reliably decreased memory precision of the initially studied item. Experiment 3 further suggested that increasing the variability of Proximal interfering items also resulted in a precision decrease. However, while these conditions included greater variability between interfering items, they also all present interference of *intermediate* similarity. Therefore, it is still unclear whether the variability of items presented during the delay period or similarity of such items to the target item is primarily responsible for precision declines. In the present and final experiment, we looked to disentangle these factors by manipulating interference similarity, while holding variability constant. To this end, we employed the same Proximal condition from Experiment 1 (120° spread, henceforth referred to as Prox-Centered), as well as a modified version of this condition where we shifted the 120° slice of color space from which the interfering colors were sampled by 60° either clockwise or counterclockwise (henceforth referred to as, Prox-Shifted; Figure 11). As a result, the interfering colors were chosen entirely from the right or left side of the studied item. If interference of intermediate similarity is a key factor that causes precision declines, then we should observe lower precision in the Prox-Shifted condition, relative to Prox-Centered and Baseline.

Method

Participants. Twenty-six individuals (18 female, 26 right-handed) with a mean age of 21.46 years (range: 17–34 years) were recruited to participate in Experiment 4. Based on the same exclusion criterion used previously (see Experiment 1 Exclusion Criteria section), we did not exclude any participants.

Design and procedure. Experiment 4 employed a fully within-subject, one-way, three-level design, with a factor of Interference Type (Baseline vs. Prox-Centered vs. Prox-Shifted). Participants completed 75 trials of Baseline, 75 trials of Prox-Centered, and 150 trials of Prox-Shifted, for a total of 300 experimental trials. We doubled the number of trials in the

---

[1] We also repeated these analyses for the Random condition. To this end, we found that misreports of the interfering colors also cannot account for precision decreases following Random interference (Appendix D–Figure A3).

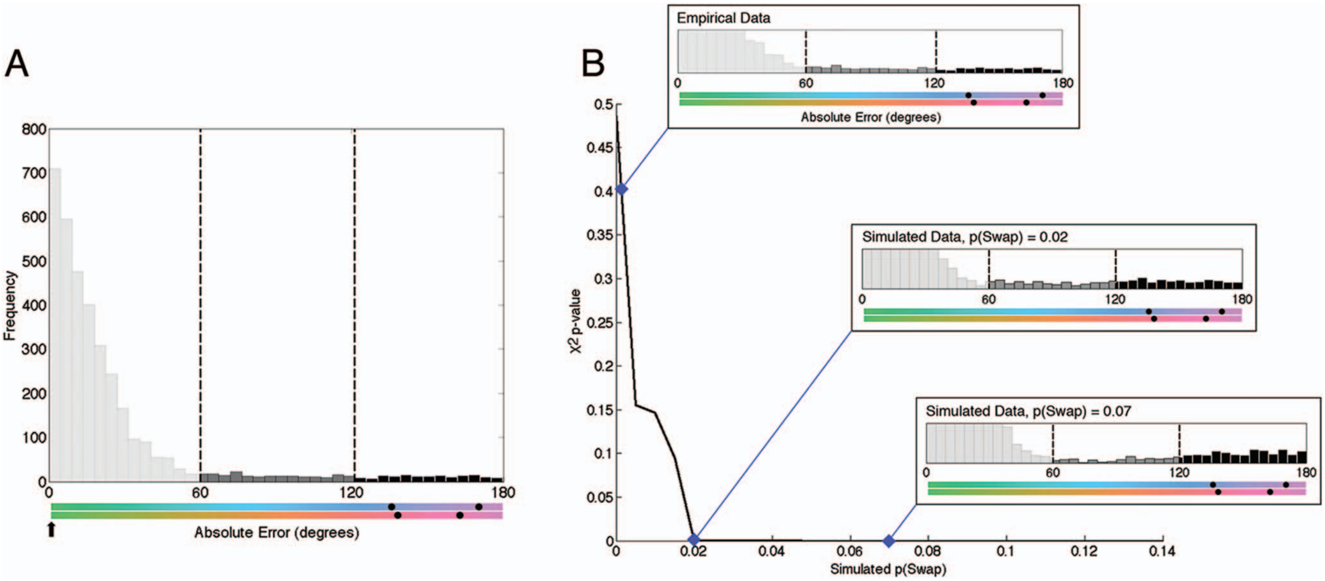

Figure 10. (A) Histogram of the absolute error in all Distal trials across all participants in Experiments 1 and 2. The leftmost region (light gray bars, 0°–59°) primarily includes responses that contribute to *p*(Mem). The central region of the figure (dark gray bars, 60°–119°) primarily represents guess responses that do not overlap with the color of Distal interfering items, while the rightmost region (black bars, 120°–179°) primarily represents guess responses that fall in the same region of color space as the Distal interfering items. An illustrative color bar below the histogram depicts an example studied color marked by the black vertical arrow at 0°, and the locations of the four unique Distal interfering colors depicted by filled black circles, for a hypothetical trial. (B) The χ^2 *p* value comparing the frequency of simulated responses between the central and rightmost histogram regions, plotted as a function of the simulated *p*(Swap) value. As the swap rate increases, the χ^2 *p* value decreases. The leftmost blue diamond (with histogram) shows the empirical data at a χ^2 *p* value of 0.4. Moving rightward, the next two blue diamonds show simulated histograms at *p*(Swap) rates of 2% (0.02) and 7% (0.07), respectively. At these simulated swap rates, χ^2 reveals highly significant differences between the central and rightmost regions of the histograms, which was not observed in the empirical data. See the online article for the color version of this figure.

Prox-Shifted condition because on half the trials, the interfering colors were shifted clockwise and on the remaining trials, they were shifted counterclockwise. We considered these together in our initial analyses, but then separated them for the subsequent analyses. Finally, we also presented 30 trials of Distal interference (identical to Experiment 1) to weaken the color association between the target and interfering items. These were intended to serve as catch trials and were not included in any analyses, as 30 trials per conditions could be insufficient to obtain reliable parameter estimates (see Experiment 1 Exclusion Criteria section). All trial types were presented in a random sequence for each participant. Furthermore, as in Experiment 1, on all trials the identity of the interfering object was the same as the studied object. The Baseline and Prox-Centered conditions were identical to their Experiment 1 counterparts. For Prox-Shifted, the slice of color space from which the interfering colors were sampled was shifted by 60° either clockwise or counterclockwise (see Figure 11). As in the previous experiments, selection of all interfering colors was constrained such that the interfering colors were separated from each other, and the studied color, by at least 5°. All other procedures were identical to the previous experiments.

Results and Discussion

Differences in *p*(Mem) and precision between conditions were tested using one-way, repeated measures ANOVAs with a factor of Interference Type (Baseline vs. Prox-Centered vs. Prox-Shifted). We found that *p*(Mem) differed between conditions, resulting in a reliable main effect of Interference Type, $F(2, 50) = 16.72$, $p < .001$; $\eta^2 = 0.40$, 95% CI [0.20, 0.62] (Figure 11A–B). Specifically, we observed an *increase* in *p*(Mem) following both Prox-Centered, $t(25) = 4.86$, $p < .001$; $d = 0.95$, 95% CI [0.57, 1.56] and Prox-Shifted, $t(25) = 2.32$, $p = .02$; $d = 0.45$, 95% CI [0.08, 0.93] interference, relative to Baseline. Furthermore, Prox-Centered increased *p*(Mem) to a greater degree than Prox-Shifted, $t(25) = 4.00$, $p < .001$; $d = 0.78$, 95% CI [0.36, 1.51]. While this may initially seem inconsistent with our previous results, we did observe a numerical increase in *p*(Mem) following Proximal interference, relative to Baseline, in both Experiments 1 and 3, which could have been attenuated by ceiling effects. Indeed, when the data are combined across these two experiments, we also find a significant increase in *p*(Mem), relative to Baseline, $t(52) = 2.18$, $p = .03$; $d = 0.30$, 95% CI [0.03, 0.58], replicating the results we observe here. Given that *p*(Mem) reflects the probability that the studied item was maintained, no matter how precisely, these results show

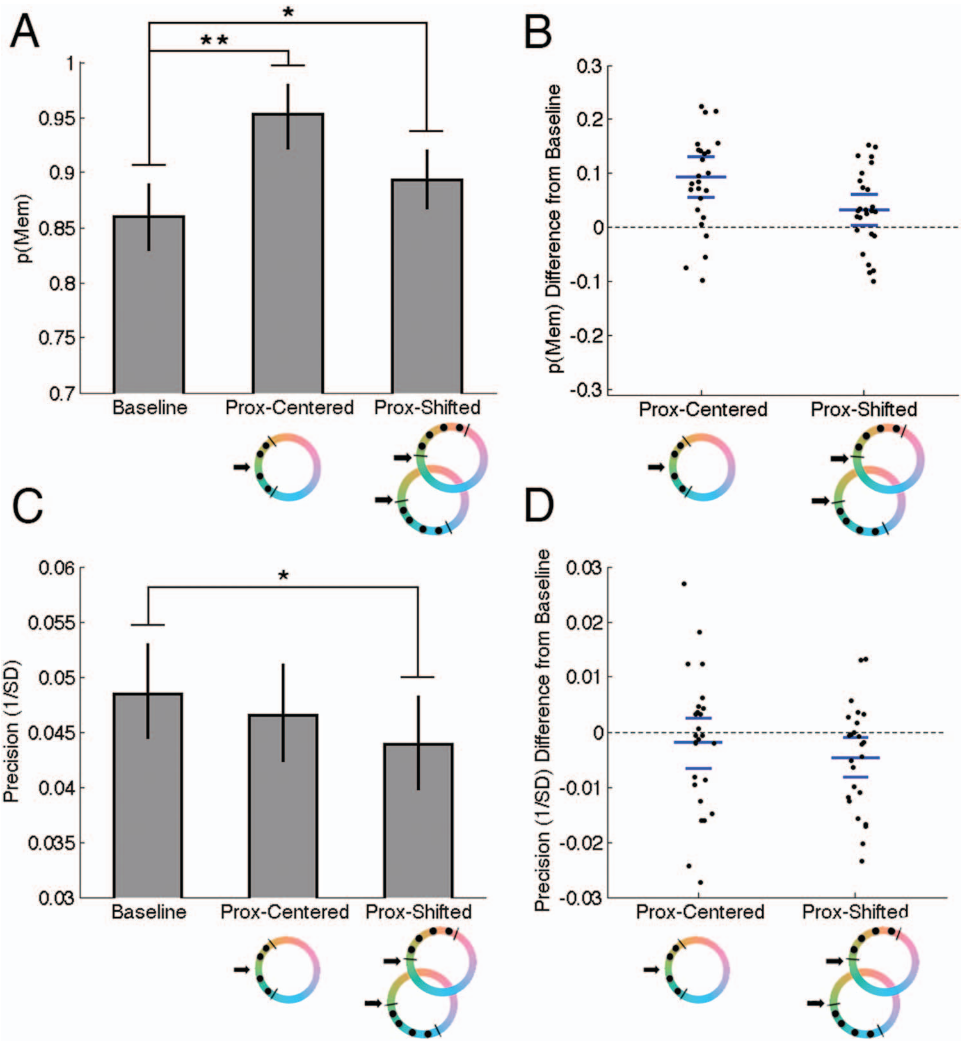

Figure 11. Results of Experiment 4. (A) p(Mem), the probability that the studied information is retained in memory, however precisely, across the interference conditions. Error bars represent 95% CIs calculated from bootstrap distributions assessing the reliability of each condition mean. (B) The interference effect in p(Mem), calculated as the difference in p(Mem) between an interference condition and the Baseline condition for individual participants. The hashed line at 0 represents the mean of the Baseline condition. The central blue line denotes the condition mean, while the outer blue lines denote the 95% bootstrap confidence interval. Any given condition is significantly different from Baseline if the hashed line at 0 falls outside of the 95% CI denoted by the outer blue lines. (C–D) Differences in memory precision between conditions, and the interference effect plotted in the same manner as (A–B), respectively. (* $p < .05$. ** $p < .01$). See the online article for the color version of this figure.

that in some cases, interference from similar information can actually increase the likelihood of recall.

However, while the benefit to p(Mem) might reflect a strengthening of the course-grained information, such a benefit might simultaneously come at a cost for the fine-grained information, measured in memory precision. For precision, we also observed a main effect of Interference Type, $F(2, 50) = 2.90$, $p = .04$; $\eta^2 = 0.10$, 95% CI [0.03, 0.31] (Figure 11C–D). Consistent with the results of Experiment 1, we did not see a decrease in precision in the Prox-Centered condition, relative to Baseline, $t(25) = 0.79$, $p = .41$; $d = 0.15$, 95% CI [−0.24, 0.59]. Interestingly, Prox-Centered resulted in the greatest benefit to p(Mem). Thus, interference from repeated presentation of highly similar colors can benefit the likelihood of recall, without showing any appreciable decrement to memory precision. However, when we increased the interference dissimilarity without changing the variability between interfering colors in the Prox-Shifted condition, we did observe a decrease in precision, relative to Baseline, $t(25) = 2.46$, $p = .01$; $d = 0.48$, 95% CI [0.11, 0.93]. Precision was also lower in Prox-Shifted, relative to Prox-Centered, $t(25) = 2.09$, $p = .04$; $d = 0.41$, 95% CI [0.03, 0.83]. These results are consistent with a surround suppression account, as memory precision was unaffected by highly similar Prox-Centered interference, but was degraded by Prox-Shifted interference of intermediate similarity.

We also asked whether declines in memory precision could be explained by an item-by-item biasing of the studied color toward the interfering colors. This was difficult to test in our previous experiments as all interfering colors were distributed symmetrically around the studied color. In the present experiment, we can examine this question by separately considering Prox-Shifted trials where the interfering colors were shifted either clockwise (henceforth, Prox-CW) or counterclockwise (henceforth, Prox-CCW). To this end, we employed a modified version of the Zhang and Luck (2008) mixture model with the mean of the Von Mises distribution as a free parameter (MemToolbox's *WithBias* function). Estimating the Von Mises mean allows the distribution of target responses to shift either clockwise or counterclockwise, thus accounting for the bias in participant responses. We fit this model separately for each participant and each of the Prox-Centered, Prox-CW, and Prox-CCW conditions. Figure 12A plots the signed bias values estimated from the fitted model. First, we do observe a bias in the studied color toward the interfering colors when they are shifted in color space. Specifically, bias increased monotonically from Prox-CCW, to Prox-Centered, to Prox-CW (Linear fit: $F(1, 25) = 80.33$, $p < .001$, $\eta^2 = 0.76$, Figure 12A, gray line). Second, it should be noted that the mean absolute bias was still substantially lower than the mean shift in the interfering colors (~8° vs. 60°; Figure 12A, outer *x*-axis bins). Next, we asked whether this bias toward the interfering colors could explain variability in memory precision in a separate condition where the interfering colors were symmetrical around the studied color, namely, Prox-Centered. To this end, we found that individuals with a greater bias toward the interfering colors (i.e., greater mean absolute bias averaged across Prox-CW and Prox-CCW) also showed lower memory precision in Prox-Centered, $r(24) = -0.42$, $p = .03$ (Figure 12B). This result is consistent with the notion that interference decreases memory precision as the representation of the studied color is sequentially biased toward each interfering color during the delay period.

General Discussion

In the present set of experiments, we found that interference from highly dissimilar information was most likely to erase the contents of memory, indexed by a decrease in the probability of recall. Furthermore, while observers were generally able to resist interference from highly similar information, we found that interference of intermediate similarity decreased memory precision (for a summary of all experiments, see Figure 13). Previous studies have extensively examined the effects of interference with respect to short- and long-term memory (e.g., Keppel & Underwood, 1962; Makovski, 2016; Sadeh et al., 2016; Watson & Lee, 2013; Wickens, 1970; Yeung et al., 2013) and immediate perception (Barense et al., 2012; Newsome et al., 2012). While the widely held similarity assumption states that interference from highly similar information should produce the greatest performance impairment, several studies have reported a memory benefit when studied items are similar to one another (Jiang et al., 2016; Johnson et al., 2009; Lin & Luck, 2009; Sims et al., 2012). Here, we take steps toward resolving this discrepancy by assuming that information could be lost from memory in at least two different ways. First, information might be erased completely, resulting in a decreased likelihood of recall. Second, memory precision might be degraded, akin to the blurring of representations that are still retained (Suchow, Brady, Fougnie, & Alvarez, 2013; Zhang & Luck, 2008). Across the first three experiments, we found that interference from highly *dissimilar* information was most likely to erase memory contents. This was measured as an increase in random guess responses, which is equivalent to a decrease in probability of recalling the studied color with any precision, *p*(Mem). Subsequent analyses showed that this

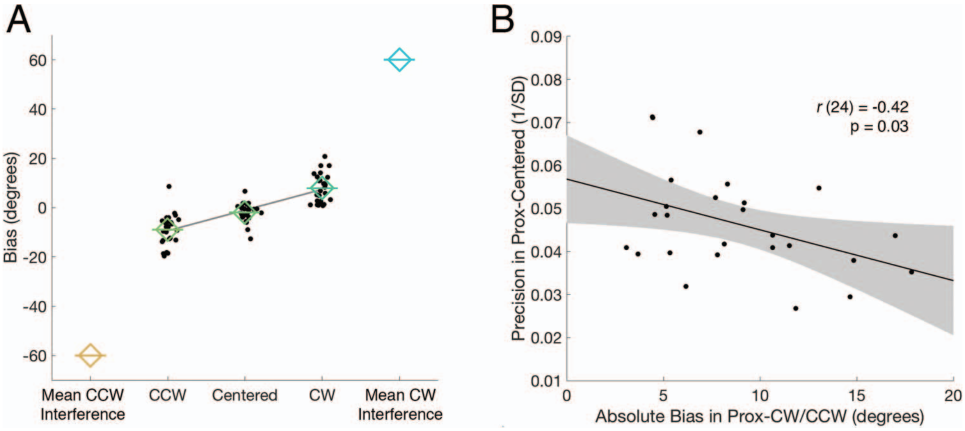

Figure 12. Analysis of bias when Prox-Shifted trials are separated into Prox-CW and CCW. (A) Each participant's bias in Prox-CCW, Prox-Centered, and Prox-CW plotted as individual dots. The gray line represents the fitted linear trend. Condition means are denoted by diamond markers. Outer diamond markers represent the mean shift of the interfering colors in Prox-CCW and CW conditions. Marker face colors illustrate the perceptual differences in color space at various degrees of bias on the *y*-axis. (B) Precision in the Prox-Centered condition plotted as a function of the absolute (i.e., unsigned) bias averaged over the Prox-CW and CCW conditions. The black line represents the fitted linear trend, and the gray shaded area represents the 95% CI of the linear fit. Participants who show a greater bias in Prox-CW/CCW tend to show lower precision in Prox-Centered. See the online article for the color version of this figure.

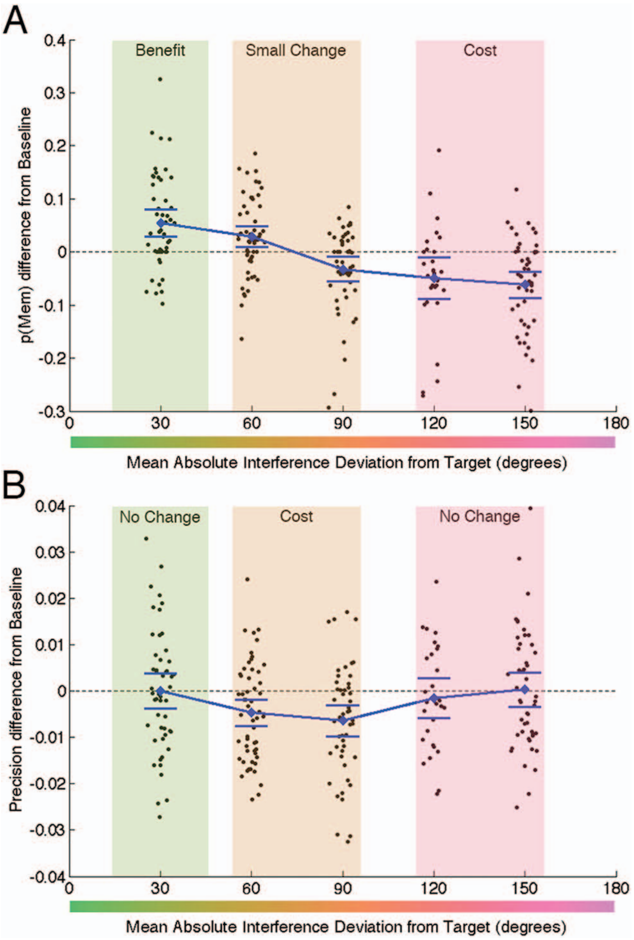

Figure 13. All data across all four experiments (E1–4) plotted as a function of the mean absolute (i.e., unsigned) distance of interfering colors from the studied color for (A) *p*(Mem) and (B) precision. Green shaded area represents highly similar interference, which benefits *p*(Mem) at no cost to precision (30°: Proximal—E1, Prox-Centered—E4). Orange shaded area represents interference of intermediate similarity, which disrupts precision, with small changes to *p*(Mem; 60°: Proximal—E3, Prox-Shifted—E4; 90°: Random—E1 and 2). Pink shaded area represents highly dissimilar interference, which decreases *p*(Mem) at no cost to precision (120°: Distal—E3; 150° Distal E1 and 2). All *x*-axis bins include two experiments of data, except for 120°, which only includes data from Distal—E3. See the online article for the color version of this figure.

decrease in *p*(Mem) was not accounted for by misreports of the interfering items (i.e., swap errors; see Between-Experiments Analyses section). Together with previous studies, these results demonstrate that the similarity assumption of interference does not apply universally, and is more nuanced than previously thought.

In the first two experiments, we also included a Random interference condition where interfering colors were broadly sampled throughout color space. In both experiments, we observed a decrease in memory precision, relative to all other conditions. However, while this condition included higher variability between interfering items, any given trial presented interfering colors that were both similar and dissimilar to memory contents (see Experiment 1 Design and Procedure section). Furthermore, while increasing the variability of Proximal interfering items also decreased precision in Experiment 3, it was unclear whether interference of high variability or intermediate similarity was the main driving factor leading to lower precision. Thus, in Experiment 4, we manipulated interference similarity while holding variability constant. Consequently, we found that while highly similar interference can be resisted, interference of intermediate similarity tended to blur memory contents by decreasing precision. This finding is consistent with a surround suppression account, which suggests that visual memories are shielded against interference from highly similar, but nonidentical information. For example, recent studies have shown that attending to a particular value in feature space (e.g., a specific shade of green) amplifies the processing of that value, while attenuating the processing of similar, but nonidentical values (Störmer & Alvarez, 2014). This effect also extends beyond attention to external features, and holds for information encoded in visual memory (Kiyonaga & Egner, 2016). Specifically, actively maintaining a color in memory can serve as an attentional template, automatically biasing attention to any external items with shared features (Olivers, 2009; Soto et al., 2008; Sun, Shen, Shaw, Cant, & Ferber, 2015). Kiyonaga and Egner (2016) found that this effect is attenuated for stimuli that were highly similar to the attentional template maintained in memory. In a subsequent experiment, they also showed that two items simultaneously held in memory could suppress each other if they were sufficiently similar. Thus, highly similar Proximal interfering colors may have fallen within the suppressive surrounding region of the target color, which attenuated their processing and protected against interference. Conversely, an item of intermediate similarity falls outside of this protective region of color space, causing interference and thus, resulting in a decrease in memory precision for the remembered item.

As previously mentioned, some studies have reported a memory benefit when nontarget items are similar to the target (Jiang et al., 2016; Johnson et al., 2009; Lin & Luck, 2009; Sims et al., 2012). Consistent with these studies, we also found that Proximal interference tended to benefit the probability of recall (i.e., increased *p*(Mem)). Also notably, we observed that these benefits can come without any appreciable cost to precision, as in Experiment 4, where we observed the largest increase in *p*(Mem) following Prox-Centered interference, which did not significantly decrease precision. Thus, in some cases, highly similar interference can reinforce the coarse-grained aspects of memory, without substantially altering the fine-grained aspects. We speculate that highly similar interference can act as a stream of imprecise reminders, reinforcing the coarse-grained aspects of representation. The effect of such imprecise reminders on fine-grained aspects of a given representation, however, may be modulated by surround suppression.

What are the mechanisms by which interference blurs or erases the contents of memory? With respect to erasure, previous computational modeling work has suggested that high featural variance (i.e., high dissimilarity) between remembered items imposes an additional tax on visual memory capacity (Jiang et al., 2016; Sims et al., 2012). Other work has suggested that information held in visual memory is susceptible to "sudden death," where items spontaneously terminate with a loss of both fine- and coarse-grained information (Zhang & Luck, 2009). Although, to our knowledge, no study has explicitly examined the relationship between featural variance and probability of sudden death, such a relationship could explain why Distal interference tends to erase memory contents. Specifically, in our experiments participants performed a one-back task on the interfering items, which requires the sequential loading of each item into memory (e.g., Fidalgo et al., 2016). Relative to the other conditions, the simultaneous retention of the studied item and any single Distal interfering item imposes greater featural variance, which could increase the probability of sudden death for the studied item.

Another interesting possibility is that the studied item is not erased, but rather, cannot be retrieved at the time of test. Indeed, retrieval failure is a well-documented phenomenon (Fougnie, Brady, & Alvarez, 2014; Rajsic, Swan, Wilson, & Pratt, 2017; Shiffrin, 1970), and it is possible that given initial retrieval failure, the studied item could nonetheless be successfully recalled in a subsequent test or with a superior memory cue. In the present experiments, there are two primary factors that might minimize the possibility of retrieval failure. First, our task did not employ time constraints during test, which would have allowed observers to utilize several retrieval attempts in a given test session. Second, given that cue-strength is a known factor in determining retrieval outcomes (Tulving, 1974), we utilized a strong retrieval cue at test in Experiment 2, where the object identity of the interfering items differed from the study item. In a recent study, Rajsic, Swan, Wilson, and Pratt (2017) demonstrated that at low set sizes, retrieval was most accurate when cues contained all features that were present at study, with the exception of the tested feature. In Experiment 2, we presented observers with the complex, real-world object they viewed during study as the retrieval cue. Critically, this cue did not overlap with the interfering object identities. Thus, in Experiment 2, we maximized the accessibility of the studied color, yet still observed a decrease in *p*(Mem) following Distal interference. However, despite our best efforts, we cannot completely rule out a retrieval failure interpretation of our current results. This interpretation would suggest that Distal interference renders memory content temporarily inaccessible at test, and interestingly, that retrieval mechanisms operate differently across interference conditions. That is, while interference of intermediate similarity blurs the precision of the studied item, highly dissimilar interference renders memories temporarily inaccessible. Future studies could examine this interesting theoretical issue by including additional nontested features at study and varying the number of features shown at test to manipulate cue strength. Crossing such a retrieval cue manipulation with different interference conditions would allow future work to determine whether retrieval failure is in fact more likely following Distal interference.

Turning now to the potential mechanisms by which memories are blurred, several studies using the memory masking paradigm have shown that presentation of a single interfering item during a retention period tends to bias memory representations toward it in feature space (Nemes et al., 2011, 2012; Van der Stigchel, Merten, Meeter, & Theeuwes, 2007; Rademaker et al., 2015). Similar findings have been reported whereby the representation of a remembered item is drawn toward the average of concurrently remembered items (Brady & Alvarez, 2011; Dubé, Zhou, Kahana, & Sekuler, 2014). In line with this idea, recently encountered information has been shown to bias judgments of immediate perception (Fischer & Whitney, 2014; Huang & Sekuler, 2010). Specifically, judgments of Gabor orientations are systematically biased toward the orientations seen on recent trials (Fischer & Whitney, 2014). Thus, in our experiments, the studied representation could have been sequentially biased toward each interfering item during the one-back task, akin to perceptual averaging. In support of this idea, we found that Proximal interfering items that were shifted either clockwise or counterclockwise tended to bias the studied representation toward them in color space (Experiment 4). Furthermore, individual differences in the degree of bias predicted memory precision in a separate condition in which the interfering items were symmetrically distributed around the studied item in color space. Specifically, individuals who showed a greater bias toward shifted interfering colors also showed lower precision when interfering colors were not shifted. Thus, we argue that when two items are considered (e.g., the studied color and a single interfering color), perceptual averaging would manifest as a biasing of the studied color toward the interfering color. When multiple interfering items are distributed symmetrically (with noise) around the studied color in color space, we argue that the same process of perceptual averaging manifests as lower precision for the studied color. Thus, we suggest that decreases in precision may result from an item-by-item biasing of the studied item toward each interfering item, ultimately resulting in a blurred memory representation.

Relatedly, decreases in precision following intermediate interference may alternatively be explained by the notion of competition (Sadil & Cowell, 2017). This interpretation suggests that the target item and the interfering items are each represented in an intact and precise manner at test. However, because all items are highly familiar, they compete for selection at retrieval. In verbal learning tasks for instance, the number of items correctly recalled from a word list decreases as competing words become more similar to one another. A common explanation of this effect is that a retrieval cue's strength becomes diluted over a greater number of items, thereby decreasing the probability of retrieval (cue overload, e.g., Wixted & Rohrer, 1993). The association of many items to a single retrieval cue may also increase the frequency of misreports, or "swaps," whereby participants report an interfering color instead of the studied color at test. As previously discussed, it is difficult both empirically and conceptually to parse an imprecise target response from a misreport of an interfering item in the Proximal condition. Empirically, as a nontarget item becomes increasingly similar to a studied item, their associated response distributions begin to overlap, eventually to an indistinguishable degree. As a result, responses from the nontarget distribution (i.e., swap errors) may manifest as decreases in precision. Conceptually, the difference between (a) representing the interfering items and the target item separately and precisely (as suggested by a competition account); and (b) retaining one imprecise target representation that is so degraded that it is reasonably well-matched by a wide range of feature values, becomes extremely difficult to disentangle when even precise representations of interfering items overlap with the target distribution. Despite the difficulties of examining misreports in our Proximal conditions, we suggest that a perceptual averaging account provides the most consistent framework for the current results in the context of prior research (Fischer & Whitney, 2014; Huang & Sekuler, 2010; Nemes et al., 2012, 2011; Rademaker et al., 2015; Van der Stigchel et al., 2007) as well as the most parsimonious account of all our results on precision. For example, we showed that misreports did not account for decreases in precision in the Random condition (see Appendix D). Similarly, Rademaker et al. (2015) found that misreports did not account for the observed 3° of bias toward an interfering Gabor when observers reported the orientation of a target Gabor. Though it is unclear whether this holds true when one or more interfering items are highly similar to the target, future studies could utilize concurrent neuroimaging to more directly assess whether there exists one blurred, or many competing, representations at test.

An alternative explanation for how Proximal interference decreased precision in our task may be via a color category-based strategy whereby participants exclusively relied on verbal labels to guide their responses (e.g., "All the items I saw on this trial were green-ish, so I'll pick something green-ish at test."). This type of strategy, whether explicit or implicit, could account for the increase in *p*(Mem) following Proximal interference. Indeed, recent work has shown that applying verbal color labels in continuous free-recall tasks can boost visual memory capacity (Souza & Skóra, 2017). In their experiments, participants studied a series of colors, and subsequently attempted to recall the exact values on a color wheel. In one condition, participants were encouraged to verbally label the colors, while in another condition, participants repeated "bababa" throughout the trial (i.e., articulatory suppression), thus preventing the use of labels. In another experiment, participants generated verbal labels and reported values on the color wheel using only the labels, with no visual information. The results showed that verbal labeling boosted memory capacity. Interestingly, verbal labels did not overwrite the continuous visual representation of the studied colors. However, the use of a verbal label alone (without visual information) resulted in a more categorical, and less continuous mode of responding. Based on these results, it is unlikely that participants in our task relied *exclusively* on a verbal strategy as categorical labels alone cannot allow for discrimination of fine-grained differences between the studied color and the highly similar interfering colors from the Proximal 120° condition (e.g., the colors from the Experiment 1 pane in Figure 7). More specifically, in Experiments 1 and 4, we found that observers could resist highly similar interference, suggesting that observers must have also maintained a continuous visual representation, which allows for fine-grained discrimination.

In our experiments, the to-be-remembered colors were presented on real-world objects (e.g., a t-shirt, a couch). In Experiment 2, we asked whether participants obligatorily encoded a conjunctive representation of color and object identity, despite the fact that object identity was task-irrelevant. Based on the representational hierarchical model (Cowell, Bussey, & Saksida, 2010; Graham, Barense, & Lee, 2010), object representations are more resistant to interference if multiple features (e.g., color and shape) are stored as a conjunction, which reduces feature-level overlap with interfering objects (Bartko et al., 2010). Thus, on half of all trials in Experiment 2, we presented the interfering colors on an object that differed from the one presented during study. We reasoned that if participants obligatorily encoded a conjunction of color and object identity, interference might be reduced when the interfering object's identity was different. However, object identity did not interact with the interference effects, suggesting that only object color was maintained. Furthermore, this finding is consistent with past visual short- and long-term memory studies reporting that object features can be stored and forgotten independently (Brady, Konkle, Alvarez, & Oliva, 2013; Fougnie & Alvarez, 2011).

One notable exception is a study by Nemes, Parry, Whitaker, and McKeefry (2012). Using a memory masking paradigm with color, these researchers found that an interfering item presented during a retention interval resulted in the greatest disruption when it was isoluminant with the studied item. Put simply, Nemes et al. (2012) found that an interfering color (i.e., hue) was less effective at disrupting memory contents when its luminance did not match the luminance of the remembered color. While this finding may initially appear inconsistent with our results, we speculate that this discrepancy could be explained by appealing to the distinction between integral and separable feature dimensions (Cant & Goodale, 2009; Cant, Large, McCall, & Goodale, 2008; Cant, Sun, & Xu, 2015; Garner, 1976). Specifically, separable feature dimensions (e.g., color and shape) can be both manipulated and perceived orthogonally (Bae & Flombaum, 2013). Conversely, while integral feature dimensions (e.g., color and luminance) can be manipulated orthogonally, they cannot be perceived independently from one another. Thus, we argue that a fruitful avenue for future research would be to increase the likelihood of encoding a conjunctive representation at study, either through the use of integral feature combinations, or by including both object identity and color as task-relevant features, when investigating how encoding and maintenance of conjunctive representations might help in resisting interference (Bae & Flombaum, 2013; Bartko et al., 2010).

In the present set of experiments, our task employed a short retention interval (~4 s), thus it may be tempting to conclude that our results pertain to visual STM. Alternatively, because we presented multiple interfering items during this interval, the studied representation may not have been maintained in an active state, and other readers might assume our results pertain to retrieval from visual long-term memory (Fidalgo et al., 2016; O'Neil et al., 2015; Watson & Lee, 2013). Creating a "process-pure" task extends far beyond the scope of the current experiments, as tasks that are traditionally thought to measure processing in one store often receive contributions from another. For instance, the complex span task, which was long regarded as the "gold standard" measure of working memory, has been shown to also rely on cue-dependent retrieval of long-term memories (Unsworth & Engle, 2007; Unsworth & Spillers, 2010). In a similar vein, while the medial temporal lobes have been thought of as a dedicated long-term memory system (Squire & Wixted, 2011), a growing body of evidence has implicated their involvement in both STM (Olson, Moore, Stark, & Chatterjee, 2006; Oztekin, Davachi, & McElree, 2010; Stern, Sherman, Kirchhoff, & Hasselmo, 2001) and perception (Barense et al., 2012; Lee et al., 2005). Given these challenges, several models of memory argue that short- and long-term memories are fundamentally integrated (cf., Cowan, 2016).

Of these models, perhaps the most relevant to the current work is the long-term working memory model of Ericsson and Kintsch (1995). This view argues that expert performance (e.g., reading, playing chess) requires active maintenance and manipulation of so much information that such tasks must depend on both temporary memory storage and efficient retrieval of stable long-term memories. In this way, long-term working memory representations are both stable, and rapidly accessible through temporarily maintained retrieval cues (Ericsson & Kintsch, 1995). For instance, skilled activities such as reading can be interrupted and resumed with minimal effort, and can be explained by long-term working memory, but not by active short-term maintenance alone. Similar to this example, our task required observers to first encode an object, direct attention toward a separate task (i.e., one-back), and then retrieve the initially studied item. Thus, it is conceivable that successful performance on our task requires both actively maintained, and more stably encoded representations. Also in line with this trend, recent neural models have argued that brain regions, such as the medial temporal lobes, can be best characterized by the type of information they represent, rather than by a specific cognitive process (e.g., perception vs. long vs. STM; Barense et al., 2012; Cowell, Bussey, & Saksida, 2010; Lee et al., 2005). Finally, others have emphasized that memory is intrinsic to neural computation in nearly all brain regions, and future research should focus on processing over multiple timescales, as opposed to just two timescales (i.e., short and long; Hasson, Chen, & Honey, 2015). Following these ideas, we do not draw strong distinctions between short- and long-term memory, as applied to our results. Rather, we utilize our experimental design as a means to characterize interference, a process known to affect and disrupt information processing over multiple timescales.

While some previous studies have reported experimental manipulations that influence both *p*(Mem) and memory precision (e.g., directed forgetting, Fawcett, Lawrence, & Taylor, 2016), others have reported manipulations that affect one parameter, while sparing the other (e.g., Brady et al., 2013; Rajsic, Sun, Huxtable, Pratt, & Ferber, 2016; Sutterer & Awh, 2016; Zhang & Luck, 2009). Indeed, these previous studies motivated our assumption that memory precision and probability of recall are, to some degree, independent of one another (Richter et al., 2016). While studies of this nature make strong contributions to our understanding of visual memory, more definitive evidence could be attained through direct comparisons between memory parameters. Comparisons between separate dependent variables can be difficult, but a principled way to do so is through Bayesian model comparison (e.g., Dowd, Kiyonaga, Beck, & Egner, 2015). In their recent work, Dowd and colleagues fit a series of 20 mixture models to data from a continuous free-recall task. The models differed with respect to (a) whether precision and *p*(Mem) were variable or constant between conditions, and (b) the direction of the difference between conditions. Using this approach, they showed that cognitive control of memory-guided attention affected the probability of recall, but not precision of the remembered item. Furthermore, by utilizing a Bayesian model comparison approach, they were able to provide direct evidence that an experimental manipulation affects one memory parameter, but not another. It should be noted that large-scale model comparisons of this nature are not standard practice, and our current experiments were not originally intended for such analyses. However, we do recognize that inference through model comparison is generally superior to frequentist statistical inference (Kruschke, 2011) and should be utilized more often in the field moving forward.

In conclusion, we argue the present results provide an important step toward a more comprehensive understanding of how interference drives the loss of information from cognitive systems. We show that highly dissimilar interference increases the likelihood of completely erasing the originally studied information. Whereas highly similar interference can be resisted, interference of intermediate similarity tends to blur memory contents by decreasing precision (see Figure 13). Our results underscore the idea that interference can disrupt the contents of memory in different ways, and should not be described along a single continuum.

References

Bae, G. Y., & Flombaum, J. I. (2013). Two items remembered as precisely as one: How integral features can improve visual working memory. *Psychological Science, 24,* 2038–2047. http://dx.doi.org/10.1177/0956797613484938

Barense, M. D., Groen, I. I., Lee, A. C., Yeung, L. K., Brady, S. M., Gregori, M., . . . Henson, R. N. (2012). Intact memory for irrelevant information impairs perception in amnesia. *Neuron, 75,* 157–167. http://dx.doi.org/10.1016/j.neuron.2012.05.014

Bartko, S. J., Cowell, R. A., Winters, B. D., Bussey, T. J., & Saksida, L. M. (2010). Heightened susceptibility to interference in an animal model of amnesia: Impairment in encoding, storage, retrieval—Or all three? *Neuropsychologia, 48,* 2987–2997. http://dx.doi.org/10.1016/j.neuropsychologia.2010.06.007

Bays, P. M. (2016). Evaluating and excluding swap errors in analogue tests of working memory. *Scientific Reports, 6,* 19203. http://dx.doi.org/10.1038/srep19203

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,* 1–11.

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

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. http://dx.doi.org/10.1177/0956797610397956

Brady, T. F., & Alvarez, G. A. (2015). Contextual effects in visual working memory reveal hierarchically structured memory representations. *Journal of Vision, 15,* 1–24.

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 of the United States of America, 105,* 14325–14329. http://dx.doi.org/10.1073/pnas.0803390105

Brady, T. F., Konkle, T., Alvarez, G. A., & Oliva, A. (2013). Real-world objects are not represented as bound units: Independent forgetting of different object details from visual memory. *Journal of Experimental Psychology: General, 142,* 791–808. http://dx.doi.org/10.1037/a0029649

Brady, T. F., Konkle, T., Gill, J., Oliva, A., & Alvarez, G. A. (2013). Visual long-term memory has the same limit on fidelity as visual working memory. *Psychological Science, 24,* 981–990. http://dx.doi.org/10.1177/0956797612465439

Brady, T. F., Störmer, V. S., & Alvarez, G. A. (2016). Working memory is not fixed-capacity: More active storage capacity for real-world objects than for simple stimuli. *Proceedings of the National Academy of Sciences of the United States of America, 113,* 7459–7464. http://dx.doi.org/10.1073/pnas.1520027113

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

Cant, J. S., & Goodale, M. A. (2009). Asymmetric interference between the perception of shape and the perception of surface properties. *Journal of Vision, 9,* 1–20.

Cant, J. S., Large, M. E., McCall, L., & Goodale, M. A. (2008). Independent processing of form, colour, and texture in object perception. *Perception, 37,* 57–78. http://dx.doi.org/10.1068/p5727

Cant, J. S., Sun, S. Z., & Xu, Y. (2015). Distinct cognitive mechanisms involved in the processing of single objects and object ensembles. *Journal of Vision, 15,* 1–21.

Cant, J. S., & Xu, Y. (2015). The impact of density and ratio on object-ensemble representation in human anterior-medial ventral visual cortex. *Cerebral Cortex, 25,* 4226–4239. http://dx.doi.org/10.1093/cercor/bhu145

Conrad, R. (1964). Acoustic confusions in immediate memory. *British Journal of Psychology, 55,* 75–84. http://dx.doi.org/10.1111/j.2044-8295.1964.tb00899.x

Cowan, N. (2016). The many faces of working memory and short-term storage. *Psychonomic Bulletin & Review*. Advance online publication. http://dx.doi.org/10.3758/s13423-016-1191-6

Cowell, R. A., Bussey, T. J., & Saksida, L. M. (2010). Components of recognition memory: Dissociable cognitive processes or just differences in representational complexity? *Hippocampus, 20,* 1245–1262. http://dx.doi.org/10.1002/hipo.20865

Dowd, E. W., Kiyonaga, A., Beck, J. M., & Egner, T. (2015). Quality and accessibility of visual working memory during cognitive control of attentional guidance: A Bayesian model comparison approach. *Visual Cognition, 23,* 337–356. http://dx.doi.org/10.1080/13506285.2014.1003631

Dubé, C., Zhou, F., Kahana, M. J., & Sekuler, R. (2014). Similarity-based distortion of visual short-term memory is due to perceptual averaging. *Vision Research, 96,* 8–16. http://dx.doi.org/10.1016/j.visres.2013.12.016

Efron, B. (1992). *Bootstrap methods: Another look at the jackknife*. New York, NY: Springer New York.

Ericsson, K. A., & Kintsch, W. (1995). Long-term working memory. *Psychological Review, 102,* 211–245. http://dx.doi.org/10.1037/0033-295X.102.2.211

Faul, F., Erdfelder, E., Buchner, A., & Lang, A.-G. (2009). Statistical power analyses using G*Power 3.1: Tests for correlation and regression analyses. *Behavior Research Methods, 41,* 1149–1160. http://dx.doi.org/10.3758/BRM.41.4.1149

Fawcett, J. M., Lawrence, M. A., & Taylor, T. L. (2016). The representational consequences of intentional forgetting: Impairments to both the probability and fidelity of long-term memory. *Journal of Experimental Psychology: General, 145,* 56–81. http://dx.doi.org/10.1037/xge0000128

Fidalgo, C. O., Changoor, A. T., Page-Gould, E., Lee, A. C., & Barense, M. D. (2016). Early cognitive decline in older adults better predicts object than scene recognition performance. *Hippocampus, 26,* 1579–1592. http://dx.doi.org/10.1002/hipo.22658

Fischer, J., & Whitney, D. (2014). Serial dependence in visual perception. *Nature Neuroscience, 17,* 738–743. http://dx.doi.org/10.1038/nn.3689

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,* 1–12.

Fougnie, D., Brady, T. F., & Alvarez, G. A. (2014, May). *If at first you don't retrieve, try, try again: The role of retrieval failures in visual working memory*. Poster presented at the 14th Annual Meeting of the Vision Sciences Society, St. Pete Beach, FL.

Fox, J., & Weisberg, S. (2011). *An {R} Companion to Applied Regression* (2nd ed.). Thousand Oaks, CA: Sage. Retrieved from http://socserv.socsci.mcmaster.ca/jfox/Books/Companion

Gao, Z., Yu, S., Zhu, C., Shui, R., Weng, X., Li, P., & Shen, M. (2016). Object-based encoding in visual working memory: Evidence from memory-driven attentional capture. *Scientific Reports, 6,* 22822. http://dx.doi.org/10.1038/srep22822

Garner, W. R. (1976). Interaction of stimulus dimensions in concepts and choice processes. *Cognitive Psychology, 8,* 98–123. http://dx.doi.org/10.1016/0010-0285(76)90006-2

Graham, K. S., Barense, M. D., & Lee, A. C. (2010). Going beyond LTM in the MTL: A synthesis of neuropsychological and neuroimaging findings on the role of the medial temporal lobe in memory and perception. *Neuropsychologia, 48,* 831–853. http://dx.doi.org/10.1016/j.neuropsychologia.2010.01.001

Hasson, U., Chen, J., & Honey, C. J. (2015). Hierarchical process memory: Memory as an integral component of information processing. *Trends in Cognitive Sciences, 19,* 304–313. http://dx.doi.org/10.1016/j.tics.2015.04.006

Huang, J., & Sekuler, R. (2010). Distortions in recall from visual memory: Two classes of attractors at work. *Journal of Vision, 10,* 1–27.

Jiang, Y. V., Lee, H. J., Asaad, A., & Remington, R. (2016). Similarity effects in visual working memory. *Psychonomic Bulletin & Review, 23,* 476–482. http://dx.doi.org/10.3758/s13423-015-0905-5

Johnson, J. S., Spencer, J. P., Luck, S. J., & Schöner, G. (2009). A dynamic neural field model of visual working memory and change detection. *Psychological Science, 20,* 568–577. http://dx.doi.org/10.1111/j.1467-9280.2009.02329.x

Kahana, M. J., & Sekuler, R. (2002). Recognizing spatial patterns: A noisy exemplar approach. *Vision Research, 42,* 2177–2192. http://dx.doi.org/10.1016/S0042-6989(02)00118-9

Keppel, G., & Underwood, B. J. (1962). Proactive inhibition in short-term retention of single items. *Journal of Verbal Learning and Verbal Behavior, 1,* 153–161. http://dx.doi.org/10.1016/S0022-5371(62)80023-1

Kiyonaga, A., & Egner, T. (2016). Center-Surround Inhibition in Working Memory. *Current Biology, 26,* 64–68. http://dx.doi.org/10.1016/j.cub.2015.11.013

Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., & Broussard, C. (2007). What's new in Psychtoolbox-3. *Perception, 36,* 1–89.

Kruschke, J. K. (2011). Bayesian assessment of null values via parameter estimation and model comparison. *Perspectives on Psychological Science, 6,* 299–312. http://dx.doi.org/10.1177/1745691611406925

Lalonde, J., & Chaudhuri, A. (2002). Task-dependent transfer of perceptual to memory representations during delayed spatial frequency discrimination. *Vision Research, 42,* 1759–1769. http://dx.doi.org/10.1016/S0042-6989(02)00111-6

Lee, A. C., Bussey, T. J., Murray, E. A., Saksida, L. M., Epstein, R. A., Kapur, N., . . . Graham, K. S. (2005). Perceptual deficits in amnesia: Challenging the medial temporal lobe 'mnemonic' view. *Neuropsychologia, 43,* 1–11. http://dx.doi.org/10.1016/j.neuropsychologia.2004.07.017

Lin, P. H., & Luck, S. J. (2009). The influence of similarity on visual working memory representations. *Visual Cognition, 17,* 356–372. http://dx.doi.org/10.1080/13506280701766313

Magnussen, S., & Dyrnes, S. (1994). High-fidelity perceptual long-term memory. *Psychological Science, 5,* 99–105. http://dx.doi.org/10.1111/j.1467-9280.1994.tb00638.x

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

Makovski, T. (2016). Does proactive interference play a significant role in visual working memory tasks? *Journal of Experimental Psychology: Learning, Memory, and Cognition*. Advance online publication.

Marshall, L., & Bays, P. M. (2013). Obligatory encoding of task-irrelevant features depletes working memory resources. *Journal of Vision, 13,* 1–13.

McKeefry, D. J., Burton, M. P., & Vakrou, C. (2007). Speed selectivity in visual short term memory for motion. *Vision Research, 47,* 2418–2425. http://dx.doi.org/10.1016/j.visres.2007.05.011

Michael, E., de Gardelle, V., & Summerfield, C. (2014). Priming by the variability of visual information. *Proceedings of the National Academy of Sciences of the United States of America, 111,* 7873–7878. http://dx.doi.org/10.1073/pnas.1308674111

Nemes, V. A., Parry, N. R. A., Whitaker, D., & McKeefry, D. J. (2012). The retention and disruption of color information in human short-term visual memory. *Journal of Vision, 12,* 1–14.

Nemes, V. A., Whitaker, D., Heron, J., & McKeefry, D. J. (2011). Multiple spatial frequency channels in human visual perceptual memory. *Vision Research, 51,* 2331–2339. http://dx.doi.org/10.1016/j.visres.2011.09.003

Newsome, R. N., Duarte, A., & Barense, M. D. (2012). Reducing perceptual interference improves visual discrimination in mild cognitive impairment: Implications for a model of perirhinal cortex function. *Hippocampus, 22,* 1990–1999. http://dx.doi.org/10.1002/hipo.22071

Olivers, C. N. (2009). What drives memory-driven attentional capture? The effects of memory type, display type, and search type. *Journal of Experimental Psychology: Human Perception and Performance, 35,* 1275–1291. http://dx.doi.org/10.1037/a0013896

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. http://dx.doi.org/10.1162/jocn.2006.18.7.1087

O'Neil, E. B., Watson, H. C., Dhillon, S., Lobaugh, N. J., & Lee, A. C. (2015). Multivariate fMRI and eye tracking reveal differential effects of visual interference on recognition memory judgments for objects and scenes. *Journal of Cognitive Neuroscience, 27,* 1708–1722. http://dx.doi.org/10.1162/jocn_a_00816

Öztekin, I., Davachi, L., & McElree, B. (2010). Are representations in working memory distinct from representations in long-term memory? Neural evidence in support of a single store. *Psychological Science, 21,* 1123–1133. http://dx.doi.org/10.1177/0956797610376651

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

Preacher, K. J., & Hayes, A. F. (2008). Asymptotic and resampling strategies for assessing and comparing indirect effects in multiple mediator models. *Behavior Research Methods, 40,* 879–891. http://dx.doi.org/10.3758/BRM.40.3.879

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

Rajsic, J., Sun, S. Z., Huxtable, L., Pratt, J., & Ferber, S. (2016). Pop-out and pop-in: Visual working memory advantages for unique items. *Psychonomic Bulletin & Review, 23,* 1787–1793. http://dx.doi.org/10.3758/s13423-016-1034-5

Rajsic, J., Swan, G., Wilson, D. E., & Pratt, J. (2017). Accessibility limits recall from visual working memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. Advance online publication. http://dx.doi.org/10.1037/xlm0000387

R Core Team. (2014). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from http://www.R-project.org/

Richter, F. R., Cooper, R. A., Bays, P. M., & Simons, J. S. (2016). Distinct neural mechanisms underlie the success, precision, and vividness of episodic memory. *eLife, 5,* e18260. http://dx.doi.org/10.7554/eLife.18260

Sadeh, T., Ozubko, J. D., Winocur, G., & Moscovitch, M. (2016). Forgetting patterns differentiate between two forms of memory representation. *Psychological Science, 27,* 810–820. http://dx.doi.org/10.1177/0956797616638307

Sadil, P. S., & Cowell, R. A. (2017). A Computational Model of Perceptual and Mnemonic Deficits in Medial Temporal Lobe Amnesia. *Journal of Cognitive Neuroscience, 29,* 1075–1088. http://dx.doi.org/10.1162/jocn_a_01106

Schneegans, S., & Bays, P. M. (2017). Neural architecture for feature binding in visual working memory. *The Journal of Neuroscience, 37,* 3913–3925. http://dx.doi.org/10.1523/JNEUROSCI.3493-16.2017

Schwarz, G. (1978). Estimating the dimension of a model. *Annals of Statistics, 6,* 461–464. http://dx.doi.org/10.1214/aos/1176344136

Shiffrin, R. M. (1970). Forgetting: Trace erosion or retrieval failure? *Science, 168,* 1601–1603. http://dx.doi.org/10.1126/science.168.3939.1601

Sims, C. R., Jacobs, R. A., & Knill, D. C. (2012). An ideal observer analysis of visual working memory. *Psychological Review, 119,* 807–830. http://dx.doi.org/10.1037/a0029856

Sneve, M. H., Sreenivasan, K. K., Alnæs, D., Endestad, T., & Magnussen, S. (2015). Short-term retention of visual information: Evidence in support of feature-based attention as an underlying mechanism. *Neuropsychologia, 66,* 1–9. http://dx.doi.org/10.1016/j.neuropsychologia.2014.11.004

Soto, D., Hodsoll, J., Rotshtein, P., & Humphreys, G. W. (2008). Automatic guidance of attention from working memory. *Trends in Cognitive Sciences, 12,* 342–348. http://dx.doi.org/10.1016/j.tics.2008.05.007

Souza, A. S., & Skóra, Z. (2017). The interplay of language and visual perception in working memory. *Cognition, 166,* 277–297. http://dx.doi.org/10.1016/j.cognition.2017.05.038

Squire, L. R., & Wixted, J. T. (2011). The cognitive neuroscience of human memory since H. M. *Annual Review of Neuroscience, 34,* 259–288. http://dx.doi.org/10.1146/annurev-neuro-061010-113720

Stern, C. E., Sherman, S. J., Kirchhoff, B. A., & Hasselmo, M. E. (2001). Medial temporal and prefrontal contributions to working memory tasks with novel and familiar stimuli. *Hippocampus, 11,* 337–346. http://dx.doi.org/10.1002/hipo.1048

Störmer, V. S., & Alvarez, G. A. (2014). Feature-based attention elicits surround suppression in feature space. *Current Biology, 24,* 1985–1988. http://dx.doi.org/10.1016/j.cub.2014.07.030

Suchow, J. W., Brady, T. F., Fougnie, D., & Alvarez, G. A. (2013). Modeling visual working memory with the MemToolbox. *Journal of Vision, 13,* 1–8.

Sun, S. Z., Shen, J., Shaw, M., Cant, J. S., & Ferber, S. (2015). Automatic capture of attention by conceptually generated working memory templates. *Attention, Perception & Psychophysics, 77,* 1841–1847. http://dx.doi.org/10.3758/s13414-015-0918-1

Sutterer, D. W., & Awh, E. (2016). Retrieval practice enhances the accessibility but not the quality of memory. *Psychonomic Bulletin & Review, 23,* 831–841. http://dx.doi.org/10.3758/s13423-015-0937-x

Tingley, D., Yamamoto, T., Hirose, K., Keele, L., & Imai, K. (2014). Mediation: R Package for Causal Mediation Analysis. *Journal of Statistical Software, 59,* 1–38. http://dx.doi.org/10.18637/jss.v059.i05

Tong, K., Ji, L., Chen, W., & Fu, X. (2015). Unstable mean context causes sensitivity loss and biased estimation of variability. *Journal of Vision, 15,* 1–12.

Tulving, E. (1974). Cue-dependent forgetting: When we forget something we once knew, it does not necessarily mean that the memory trace has been lost; it may only be inaccessible. *American Scientist, 62,* 74–82.

Unsworth, N., & Engle, R. W. (2007). The nature of individual differences in working memory capacity: Active maintenance in primary memory and controlled search from secondary memory. *Psychological Review, 114,* 104–132. http://dx.doi.org/10.1037/0033-295X.114.1.104

Unsworth, N., & Spillers, G. J. (2010). Working memory capacity: Attention control, secondary memory, or both? A direct test of the dual-component model. *Journal of Memory and Language, 62,* 392–406. http://dx.doi.org/10.1016/j.jml.2010.02.001

Van der Stigchel, S., Merten, H., Meeter, M., & Theeuwes, J. (2007). The effects of a task-irrelevant visual event on spatial working memory. *Psychonomic Bulletin & Review, 14,* 1066–1071. http://dx.doi.org/10.3758/BF03193092

Viswanathan, S., Perl, D. R., Visscher, K. M., Kahana, M. J., & Sekuler, R. (2010). Homogeneity computation: How interitem similarity in visual short-term memory alters recognition. *Psychonomic Bulletin & Review, 17,* 59–65. http://dx.doi.org/10.3758/PBR.17.1.59

Watson, H. C., & Lee, A. C. (2013). The perirhinal cortex and recognition memory interference. *The Journal of Neuroscience, 33,* 4192–4200. http://dx.doi.org/10.1523/JNEUROSCI.2075-12.2013

Wickens, D. D. (1970). Encoding categories of words: An empirical approach to meaning. *Psychological Review, 77,* 1–15. http://dx.doi.org/10.1037/h0028569

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

Wixted, J. T. (2004). The psychology and neuroscience of forgetting. *Annual Review of Psychology, 55,* 235–269. http://dx.doi.org/10.1146/annurev.psych.55.090902.141555

Wixted, J. T., & Rohrer, D. (1993). Proactive interference and the dynamics of free recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 19,* 1024–1039. http://dx.doi.org/10.1037/0278-7393.19.5.1024

Yeung, L. K., Ryan, J. D., Cowell, R. A., & Barense, M. D. (2013). Recognition memory impairments caused by false recognition of novel objects. *Journal of Experimental Psychology: General, 142,* 1384–1397. http://dx.doi.org/10.1037/a0034021

Yotsumoto, Y., Kahana, M. J., Wilson, H. R., & Sekuler, R. (2007). Recognition memory for realistic synthetic faces. *Memory & Cognition, 35,* 1233–1244. http://dx.doi.org/10.3758/BF03193597

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

Appendix A

Simulations for Determining Exclusion Criteria

For all experiments, participants performed a one-back task during the interference phase to ensure that the interfering items were attended to. We performed all analyses on trials where participants correctly identified the repeating item, thus effectively holding one-back accuracy at 100% across all conditions. However, three participants from both Experiments 1 and 2 were excluded due to exceptionally low hit rates on the one-back task. Our exclusion criterion was determined through the following simulation, described below (Figure A1).

Simulation Results

Estimated parameter value plotted against the number of trials included in the simulation. For each bin on the X-axis (in steps of five), 25 data sets were simulated. The central blue line represents the mean parameter value of the 25 data sets, while the shaded blue region represents the 95% CI. The red line denotes the "ground truth" parameter value.

For each bin on the X-axis (in steps of five), we simulated a sample of 25 data sets, where each dataset included X number of trials (Figure A1). The horizontal red line denotes the "ground truth" parameter value. The true *p*(Guess) value (0.08) and *SD* value (19°) are based on the empirical parameter estimates obtained in Experiment 1, collapsed over all participants and all conditions, before exclusions. The simulation reveals that parameter estimates are quite reliable, even if each dataset includes a relatively low number of trials. However, the variability between simulated data sets appears to reach stability, never significantly deviating from the ground truth, at about 40 trials. Thus for all experiments, our analysis included all participants with at least 40 trials retained in each condition.

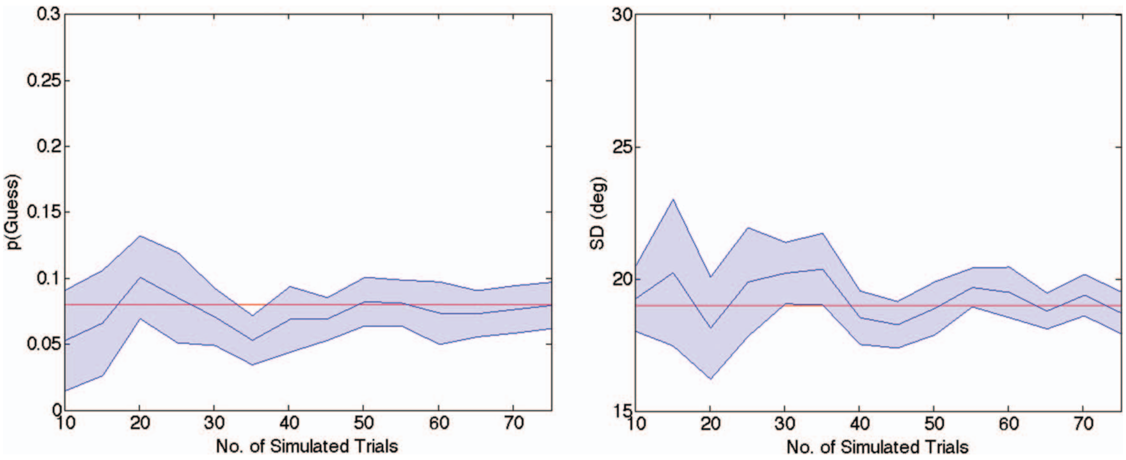

Figure A1. Simulation results. Estimated parameter value plotted against the number of trials included in the simulation. For each bin on the X-axis (in steps of 5), 25 datasets were simulated. The central blue line represents the mean parameter value of the 25 datasets, while the shaded blue region represents the 95% CI. The red line denotes the "ground truth" parameter value. See the online article for the color version of this figure.

(*Appendices continue*)

Appendix B

Mean Hit Rates for the One-Back Task for All Experiments After Exclusions (SD in Brackets)

	Baseline	Random	Distal	Proximal
Experiment 1	.76 (.12)	.87 (.09)	.79 (.12)	.80 (.13)
Experiment 2	.70 (.16)	.80 (.17)	.73 (.17)	
Experiment 3	.75 (.11)		.85 (.07)	.85 (.09)
	Baseline	Prox-Centered	Prox-CW	Prox-CCW
Experiment 4	.80 (.10)	.86 (.12)	.85 (.12)	.86 (.10)

Appendix C

Simulations of *p*(Swap) Estimation with Non-Independently Selected Interfering Items

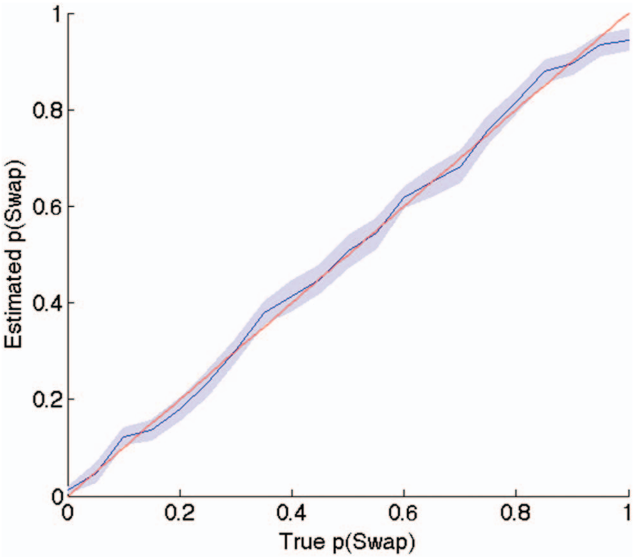

Figure A2. Simulations assessing whether *p*(Swap) parameter values could be correctly recovered when interfering items are not independently sampled in color space. Mean *p*(Swap) values for each bin plotted as the blue line, while the shaded region represents the 95% CI. All simulations use Distal interfering items drawn from a 120° slice of color space centered 180° away from the target color. All other simulation details identical to Appendix A. See the online article for the color version of this figure.

(*Appendices continue*)

Appendix D

Precision Decrease Following Random Interference is Not Accounted For By Swap Errors

In the Between-Experiments Analyses section, we showed that misreports of the interfering items (i.e., swap errors) could not account for memory erasure following Distal interference. Here, we applied the same analyses to assess whether misreports could account for decreases in precision following Random interference. Figure A3A plots Bayesian information criterion (BIC; Schwarz, 1978) differences comparing the fit of the *StandardMixtureModel* and *SwapModel* to data from the Random condition for Experiments 1 and 2. This analysis revealed that the *StandardMixtureModel* was preferred in 44 out of 49 participants.

Next, we excluded all participants that committed swap errors on more than 1% of trials (23 participants retained). For the remaining participants (bolded dots in Figure A3A), the estimated *p*(Swap) was 0.004 and the *StandardMixtureModel* was roughly six times more likely than the *SwapModel*. In this subset of participants, *p*(Mem) did not differ from Baseline following Random interference, $t(22) = 0.80$, $p = .41$; $d = 0.17$, 95% CI [−0.33, 0.49]; Figure A3B. However, we did observe a robust decrease in memory precision following Random interference, relative to Baseline, $t(22) = 4.31$, $p < .001$; $d = 0.90$, 95% CI [0.48, 1.54]; Figure A3C. Therefore, we conclude that misreports of the interfering items cannot account for decreases in memory precision following Random interference (Figure A3).

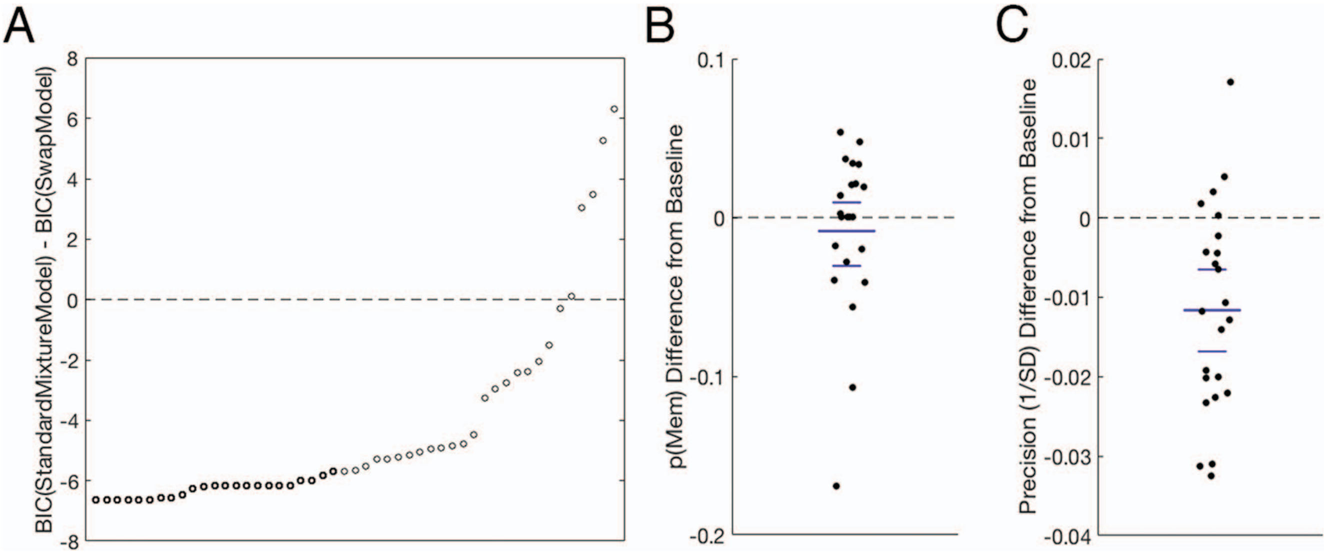

Figure A3. (A) Bayesian information criterion (BIC) differences comparing the 2-component *StandardMixtureModel* with the 3-component *SwapModel* for the Random condition across Experiments 1 and 2. Dots represent individual participants. Values represent the likelihood of a model, given the data, and competing models. Positive values indicate support for the *SwapModel*. Bolded dots represent a subsample of participants that commit swap errors on less than 1% of trials. (B & C) Analysis of memory parameters in the subsample of participants with <1% swap errors. Dots represent the difference in *p*(Mem; B), or precision (C) between Random and Baseline. The central blue line denotes the mean, while the outer blue lines denote the 95% bootstrap confidence interval. The hashed line at zero denotes the mean of the Baseline condition. See the online article for the color version of this figure.

Received September 19, 2016
Revision received July 4, 2017
Accepted July 10, 2017 ■