

Adaptive Repulsion of Long-Term Memory Representations Is Triggered by Event Similarity

🕕 😳

Avi J. H. Chanales¹, Alexandra G. Tremblay-McGaw², Maxwell L. Drascher², and Brice A. Kuhl^{2,3}

¹Department of Psychology, New York University; ²Department of Psychology, University of Oregon; and ³Institute of Neuroscience, University of Oregon

Psychological Science 2021, Vol. 32(5) 705–720 © The Author(s) 2021 Article reuse guidelines: sagepub.com/journals-permissions DOI: 10.1177/0956797620972490 www.psychologicalscience.org/PS



Abstract

We tested whether similarity between events triggers adaptive biases in how those events are remembered. We generated pairs of competing objects that were identical except in color and varied the degree of color similarity for the competing objects. Subjects (N = 123 across four experiments) repeatedly studied and were tested on associations between each of these objects and corresponding faces. As expected, high color similarity between competing objects created memory interference for object–face associations. Strikingly, high color similarity also resulted in a systematic bias in how the objects themselves were remembered: Competing objects with highly similar colors were remembered as being further apart (in color space) than they actually were. This repulsion of color memories increased with learning and served a clear adaptive purpose: Greater repulsion was associated with lower associative-memory interference. These findings reveal that similarity between events triggers adaptive-memory distortions that minimize interference.

Keywords

episodic memory, long-term memory, forgetting, open data, open materials

Received 12/17/19; Revision accepted 9/28/20

Much of our day-to-day experience is composed of similar events, and this similarity between events can lead to interference and forgetting. Thus, resolving interference is one of the most important and fundamental challenges for the memory system. Insight into how the memory system resolves interference can potentially be gleaned from the computations performed by the hippocampus—a neural structure thought to play a critical role in separating similar memories (Bakker et al., 2008; Colgin et al., 2008; Yassa & Stark, 2011). Traditionally, the hippocampus has been thought to minimize interference by forming independent representations of events, even when those events are highly similar. Recently, however, several neuroimaging studies have found that the hippocampus goes one step further: that highly similar events may, in fact, have less correlated patterns of hippocampal activity than unrelated events (Ballard et al., 2019; Chanales et al., 2017; Dimsdale-Zucker et al., 2018; Favila et al., 2016; Hulbert & Norman, 2015; Schlichting et al., 2015). These results

suggest a dependence between hippocampal representations of similar events whereby similarity triggers a repulsion such that representations of highly similar events systematically move away from each other, much like equal poles of a magnet do when they get too close. Importantly, hippocampal repulsion has also been associated with reduced memory interference (Favila et al., 2016; Hulbert & Norman, 2015), consistent with the idea that this process serves an adaptive purpose. At present, however, it is unclear whether a similar repulsion also applies to behavioral expressions of memory. When events are highly similar, are the features of those events systematically remembered as further apart than they actually were? If so, does this repulsion protect against interference?

Corresponding Author:

Brice A. Kuhl, University of Oregon, Department of Psychology E-mail: bkuhl@oregon.edu

The potential for similarity to trigger repulsion of feature memory-and the consequences of such repulsioncan be understood in the context of classical memoryinterference paradigms. For example, in the canonical A-B, A-C interference paradigm (Barnes & Underwood, 1959), a single memory cue (A) is paired with two different associates (B, C). Interference between the B and C terms is a direct consequence of the overlapping cue (A). In reality, however, cues are often partially overlapping (A_1) and A_2), and the amount of interference is proportional to the degree of overlap (Mensink & Raaijmakers, 1988; Osgood, 1949). For example, two nearly identical cars (highly similar cues) may belong to two different colleagues (associates), creating strong potential for memory interference. In contrast, if two very different cars belong to two other colleagues, the potential for memory interference will be lower. Although there is a rich literature addressing the factors and contexts that promote interference-related forgetting, these studies overwhelmingly focus on whether or not a given cue (A) is effective in eliciting retrieval of a particular associate or associates (B, C; Anderson, 2003; Mensink & Raaijmakers, 1988; Osgood, 1949). In contrast, the phenomenon of hippocampal repulsion described above motivates a different question: Does similarity among memory cues $(A_1 \text{ vs. } A_2)$ trigger a systematic exaggeration in how these cues are remembered? Given the well-established relationship between cue overlap and interference, the potential benefit of exaggerating the differences between cues is clear: It should reduce interference when corresponding associates are recalled (O'Reilly, 2004).

Here, we report a series of behavioral experiments that used a paired-associates learning paradigm (A_1-B) , A_2 -C) to test whether similarity among memory cues triggers an adaptive exaggeration in how those cues are remembered. In each of four experiments, subjects studied object-face associations with interference induced by including pairs of overlapping objectsobjects that were identical except for their color values. For example, a blue backpack and a purple backpack $(A_1 \text{ and } A_2)$ would each be associated with a distinct face (B or C; Fig. 1a). Importantly, we also parametrically varied the degree of color similarity between the overlapping object cues (i.e., A₁-A₂ similarity). On the basis of the robust memory-interference literature, we anticipated that greater color similarity between overlapping objects would result in greater interference when the face associates were recalled (Mensink & Raaijmakers, 1988; Osgood, 1949). However, the critical focus of our experiments was whether similarity between the overlapping objects would lead to distortions or bias in how the object cues themselves were remembered. Specifically, we probed memory for each object's color in order to determine whether subjects

Chanales et al.

Statement of Relevance

Experiencing similar events (e.g., two summer vacations in the same location with the same family members) creates the potential for memory confusion and errors. Given how much similarity there is among our day-to-day experiences, one of the critical functions of the memory system is to help minimize the potential for similarity-based confusion and errors. Here, we show that one of the ways the memory system accomplishes this is by exaggerating differences among highly similar events. Although these exaggerations reflect a distortion of reality, we show that this distortion is adaptive: When events are remembered as being less similar, there is less potential to confuse them in memory.

exaggerated the distance in color space between similar objects.

Our predictions were directly motivated by properties of hippocampal repulsion. Our central prediction was that subjects would exaggerate the distance between overlapping objects in color memory but only when the objects had highly similar colors. Put another way, we expected that similarity would trigger colormemory repulsion (Chanales et al., 2017; Favila et al., 2016; Hulbert & Norman, 2015; Schapiro et al., 2012). As with hippocampal repulsion, we expected that repulsion of color memories would occur gradually and as a result of experience (Chanales et al., 2017; Favila et al., 2016). Thus, we used an overtraining paradigm with the critical test of color-memory repulsion occurring at the end of the training. Our second main prediction was that repulsion would be adaptive (Colgin et al., 2008; Favila et al., 2016; Hulbert & Norman, 2015): that greater repulsion of color memories for overlapping objects would reduce interference when corresponding associates (faces) were recalled. This prediction is notable-and counterintuitive-because it posits a distortion of memories, caused by interference, that benefits memory retrieval.

Method

Stimuli, data, and relevant code are available on OSF at https://osf.io/d4pmg/.

Subjects

For all experiments, target sample sizes were identified in advance, but final sample sizes were determined by exclusion criteria. Because we did not have any way to

sample: 36 subjects were included in analyses (13 men; 18–22 years old). For Experiments 3, 38 subjects were included in analyses (6 men; 18–34 years old), and for Experiment 4, 26 subjects were included in analyses (1 man; 18–25 years old). The rationale for the smaller sample in Experiment 4 was that pilot data indicated that the critical manipulation in Experiment 4 was quite



Fig. 1. (continued on next page)

Fig. 1. Experimental design. In each of four experiments, subjects learned objectface associations (a) that contained pairs of competing objects (object images that were identical except for their color values). The similarity (color distance) between competing objects was parametrically manipulated within and across experiments. In Experiments 1, 2, and 4, there were three similarity conditions: low similarity (72° apart), moderate similarity (48° apart), and high similarity (24° apart). In Experiment 3, the conditions were moderate similarity (48° apart), high similarity (24° apart), and ultra similarity (6° apart; not shown). Each experiment began with eight training rounds (b). Each training round contained a study phase, a color-memory test, and an associative-memory test. During study, subjects viewed each object-face pair. During color-memory tests, subjects were shown a face and a gray-scale version of the associated object. Using a continuous color wheel, subjects selected (recalled) the color of the object. During associativememory tests (Experiments 1-3 only), an object image was presented, and subjects selected the associated face from a set of four options. The four face options always included the correct face (target) and the face that had been paired with the competing object (competitor). Procedures for Experiment 4 are described in Figure 5. For all experiments, subjects completed a posttraining test (c) that probed only color memory. The procedure was identical to that of the color-memory tests from the training rounds. The critical performance measure in the posttraining test was the percentage of color-memory responses that were biased away from the color of the competing object.

powerful and consistent across subjects. Exclusion criteria are described in the Procedure section. All subjects were right-handed and reported normal or correctedto-normal vision. Informed consent was obtained in accordance with procedures approved by the University of Oregon Institutional Review Board.

Materials

For all experiments, stimuli consisted of 18 object images and 36 face images (all images were 250 × 250 pixels). The object images were selected from a set of images designed to be color-rotated (Brady et al., 2013). Objects were chosen that had no strong inherent associations with particular colors. To alter the color of each object, we rotated the hue of an image through a circular color space ranging from 0° to 360°. Colors were altered by independently rotating every pixel through an equiluminant circle in Commission Internationale de l'Éclairage (CIE) 1976 L*a*b* color space. Face images were pictures of nonfamous white males gathered from the Internet. For each subject, six object images were randomly assigned to each of three colorsimilarity conditions. For Experiments 1, 2, and 4, these conditions were high similarity (24°), moderate similarity (48°), and low similarity (72°; Fig. 1a). For Experiment 3, these conditions were ultra similarity (6°) , high similarity (24°), and moderate similarity (48°). Each object was then assigned a pair of colors separated by the hue-angle-degree difference of their respective condition. For Experiments 1, 2, and 4, this was accomplished by randomly selecting 45 colors from the color space, each separated by 8° (45 × 8 = 360°). This set of 45 colors represented the set of possible colors for each subject, but only 36 of these colors were actually used (18 objects \times 2 colors per object). For Experiment 3, the ultra-similarity condition necessitated a slight modification to the color-assignment procedure: 60 colors (instead of 45) were randomly selected from the color space, each separated by 6° (instead of 8°). This set of 60 colors represented the set of possible colors for each subject, but again, only 36 of these colors were used. For all experiments, the 36 colors were then assigned to objects, according to their similarity condition, without replacement (i.e., each color was assigned only to one object). One constraint on this assignment was that, for each condition, there was an even representation across each third of the color space (1°–120°, 121°–240°, 241°–360°).

Procedure

The first part of each experiment was a series of training rounds (eight total) in which subjects learned and were tested on all of the object-face associations (Fig. 1b). Each training round was composed of a study phase, a color-memory test, and an associative-memory test. During the study phase, subjects were shown (and asked to encode) each object-face association. During the color-memory test, subjects were shown a face image and a gray-scale version of the object that was associated with that face. Using a continuous report scale (Brady et al., 2013), subjects selected the color of the object. No feedback was provided during the colormemory tests. During the associative-memory test in Experiments 1 through 3, subjects were shown an object image along with four face images, and they attempted to select the face that had been paired with that object. Importantly, the set of four faces included the face that was paired with the object (the target), the face that was paired with the competing object (the competitor), and two faces that had been paired with different objects

(noncompetitors). Feedback was provided after each response during the associative-memory test.

After the eight training rounds, subjects completed a color-memory posttraining test that repeatedly tested their memories for each object's color (Fig. 1c), again without feedback. This posttraining test was identical in format to the color-memory tests during the training rounds, but it served as a critical measure of the end point of subjects' learning. Prior to the start of each experiment, subjects were informed that every face would be paired with exactly one object but that object categories would be associated with more than one face; subjects were explicitly made aware that when object categories were associated with more than one face, the object colors would always be different (even if very similar).

Experiment 1. Experiment 1 consisted of eight training rounds and two posttraining tests. Each training round included a study phase, a color-memory test, and an associative-memory test (in that order). In each study round, subjects viewed the same 36 object–face associations. For each trial, a face and corresponding object image appeared on a white screen for 2.5 s, with the face image to left of the object. There was a 1-s intertrial interval during which a blank white screen was presented. Each object–face association was presented once per study round.

On each trial of the color-memory test, a studied face was presented to the left of its paired object, and subjects used a color wheel (Brady et al., 2013) to select the remembered color of the object. The object image initially appeared in gray scale; when subjects moved the mouse cursor, the object would appear in color. The hue was determined by the angle between the cursor location and the center of the object image. A line marker along a ring surrounding the object image indicated the current hue angle. Once subjects rotated to the desired color, they clicked the mouse to finalize their choice. There was no time limit for these responses. The color wheel was randomly rotated across trials so there was no correspondence across trials between spatial position and color.

The associative-memory tests probed memory for each object–face association. On each trial, a colored object was presented at the top of the screen, and four face images were presented beneath. Subjects were instructed to select the face that had been studied with the object image. The target face was always presented along with a face that was paired with the competing object. Thus, subjects had to discriminate between the objects' colors (in memory) in order to select the correct face. The other two face images were randomly selected and had been paired with other, noncompeting objects (noncompetitor faces). Each face served as a noncompetitor foil on two trials. Subjects made a selection using a mouse with no time limit to respond. They then indicated confidence in their response by clicking either a "sure" or an "unsure" button (these confidence ratings are not considered here). A feedback screen was then presented for 1.25 s; the feedback screen indicated whether the selected face was correct or not and also displayed the correct object–face pair.

Following the eight training rounds, subjects completed an immediate posttraining test (Day 1 posttraining test) and then returned the following day (~24 hr later) for a second posttraining test (Day 2 posttraining test). The posttraining tests were identical to the colormemory tests in the training rounds except that in the posttraining tests, each object was tested 5 times. The 180 trials of the posttraining test were divided into five blocks in which each object was tested once. The order of the trials within a block was randomized, with the constraint that an object and its competitor were not tested on successive trials. After Trial 60 (one third of trials competed) and Trial 120 (two thirds of trials completed), a screen prompted subjects to "Take a quick break" and to press a key to continue (this was done to minimize fatigue). The Day 2 posttraining test was identical to the Day 1 posttraining test except that the order of trials was rerandomized.

Experiment 2. Experiment 2 was identical to Experiment 1 except for a few procedural changes. Because of the strict performance-based exclusion criteria in Experiment 1 and the time-limit cutoff (1.5 hr), there was a high overall exclusion rate (45.2%) and a relatively small final sample of subjects (N = 23). Thus, the goal for Experiment 2 was to replicate the results from Experiment 1 but with a larger final sample (the target was a 50% increase) and a lower overall exclusion rate. We retained the same exclusion criteria but shortened the experiment to reduce the number of subjects who failed to complete the experiment in the allotted time (1.5 hr). We opted to shorten the experiment rather than extend the time limit out of concern for subject fatigue. Subject fatigue was of particular concern given that the most critical data from the entire experiment came from the posttraining test (i.e., the last round of the experiment). To shorten the experiment, we reduced the number of color-memory test rounds during training so that they occurred only during Rounds 1, 3, 5, and 7, and we imposed a time limit of 10 s on each trial of the associative-memory test and colormemory test (both during the training rounds and the posttraining test). The number of trials on the posttraining test that were excluded because of the time limit in Experiment 2 was very small (across-subjects M = 0.78%; maximum = 3.3%). Additionally, because we observed qualitatively identical results across the Day 1 and Day 2 posttraining tests in Experiment 1, we did not include the Day 2 posttraining test in Experiment 2.

Experiment 3. Experiment 3 was identical in procedure to Experiment 2 except that color-memory tests were included in each of the eight training rounds (as in Experiment 1), and the time limit for the entire experiment was extended to 2 hr. The rationale for reverting to color-memory tests every round was that the magnitude of the repulsion effect in the high-similarity condition was somewhat lower in Experiment 2 (M = 54.63%) than in Experiment 1 (M = 60.80%), and we speculated that this difference might be partly attributable to the reduction in the number of color-memory tests during training in Experiment 2. The rationale for extending the time limit was to reduce the number of subjects excluded for not completing the experiment in the allotted time.

Experiment 4. Experiment 4 differed in procedure from Experiments 1, 2, and 3, most significantly replacing the associative-memory test during the training rounds with an inference test that assessed generalization across object-face pairs. Specifically, on each trial in the inference test, one of the 36 face images appeared at the top of the screen, and four face images appeared beneath. Subjects were instructed to select which of the four face options was associated with the same object category (e.g., backpack) as the probe face at the top of the screen. Subjects made their selection using a mouse. The set of four face options always included the target face (correct response) and three other studied nontarget faces. (Note that here the competitor face was the target face.) Each face was tested once per round (i.e., 36 trials per round), and each face served as a nontarget face option on exactly three trials. Feedback was provided on each trial to indicate whether the selected face was correct; the correct face-face pairing was displayed on screen for 1 s. Otherwise, all procedures were identical to those in Experiment 3 (for details of color-similarity conditions, see the Materials section).

Exclusion criteria. Across Experiments 1 through 4, a total of 19, 18, 9, and 6 subjects, respectively, were excluded from analyses. Subjects were excluded if they failed to complete the experiment in the allotted time (1.5 hr for Experiments 1 and 2 and 2 hr for Experiments 3 and 4). For Experiments 1 through 4, this resulted in exclusion of 12, 0, 1, and 3 subjects, respectively. Additionally, subjects were excluded if they did not satisfy each of two performance-based criteria. The first performance criterion was that across the last two rounds of the associativememory test (see the Procedure section for details), noncompetitor face images were selected on no more than 2% of trials. Importantly, this exclusion criterion was orthogonal to subjects' ability to discriminate between similar colors because it required only that subjects had narrowed down the options to either the target or the competitor face. This exclusion criterion therefore specifically ensured that subjects had learned that two different faces were paired with a common object category (e.g., backpack). Across Experiments 1 through 4, a total of 6, 15, 9, and 0 subjects, respectively, failed to meet this criterion.

The second performance criterion was that the percentage of trials on the posttraining test with reaction times less than 500 ms could not exceed 15%. Given that the posttraining test required clicking a mouse and dragging a cursor along a color wheel, responses that were made in less than 500 ms were considered to be evidence of subjects rushing through the experimentwhich was a particular concern given its repetitive and tedious nature. Across Experiments 1 through 4, a total of 5, 3, 1, and 6 subjects, respectively, failed to meet this criterion. (Some subjects failed to satisfy both of the performance-based criteria-4, 5, 3, and 0 subjects, respectively, in Experiments 1-4.) It is important to emphasize that these performance-based exclusion criteria were established in advance, they were orthogonal to our effects of interest (repulsion of color memory), and they were applied uniformly across all experiments.

Measuring color memory. For color-memory tests (during training rounds and posttraining test), responses were recorded as integer values (0°-359°) reflecting hue angle on the color wheel. Although the procedure was identical for the color-memory tests during the training rounds and posttraining test, we focus for narrative clarity on different measures during each phase. During the training rounds, we focus on color error as a general measure of training-related improvement in color memory. Color error was computed as the absolute value of the hueangle difference between a subject's color response and the true color. During the posttraining test, however, because we were critically interested in whether colormemory responses exhibited bias, we focused on the percentage of responses away from the competing object's color ("away" responses). To illustrate how this measure was computed, consider the location of the target object's color value as being at 0° on the color wheel and the location of a competing object's color value at 24°. In this scenario, any responses between 181° and 359° would be counted as away from the competing color. Notably, the definition of "away" responses would be identical if the competing object color were at 6°, 48°, or 72°. For each subject and each condition, we computed the percentage of trials that fell in the "away" bin. We used this measure as opposed to mean signed error because mean signed error is highly susceptible to influence from extreme responses, whereas the percentage of "away" responses is not. Although we focus on the colorerror measure during training rounds and the percentage of "away" responses during the posttraining test, we report the percentage of "away" responses during training and the mean color error during the posttraining test



Fig. 2. Results from Experiment 1 (top row) and Experiment 2 (bottom row). Mean color-memory error (a; absolute distance between reported and target color values from the color-memory tests) and accuracy on the associative-memory test (b; percentage of trials for which the target face was selected; chance = 25%) are shown as a function of training round and similarity condition. The percentage of responses away from the competitor on the posttraining test (c) is shown as a function of training day and similarity condition. Large dots represent means, and small dots represent data from individual subjects. Symbols represent a marginally significant difference ($^{\dagger}p < .10$) and significant differences ($^{*}p < .05$, $^{**}p < .001$) from 50%. Error bars represent standard errors of the mean.

in the Supplemental Material available online (see Figs. S1 and S3, respectively).

Results

Experiment 1

Across the training rounds in Experiment 1, there were marked reductions in error on the color-memory tests, F(1, 22) = 166.2, p < .001, $\eta_G^2 = .77$ (Fig. 2a) and increases in accuracy on the associative-memory tests, F(1, 22) = 435.4, p < .001, $\eta_G^2 = .8$ (Fig. 2b; see Figs. S1 and S2 in the Supplemental Material for additional data from the training rounds). Critically, accuracy on the associative-memory

tests was strongly influenced by color similarity, F(2, 44) = 13.04, p < .001, $\eta_G^2 = .23$ (Fig. 2b). In particular, subjects were much more likely to select the face associated with the competing object (hereafter, *interference errors*) when color similarity was high (see Fig. S2), confirming that our interference manipulation was successful.

To investigate repulsion effects in color memory, we focused on the posttraining tests. We predicted that repulsion would specifically occur when competition was high (i.e., in the high-similarity condition). The critical dependent measure was the percentage of trials in each similarity condition for which subjects reported a color that was away from the color of the competing object (measures of unsigned color error are reported in Fig. S3). For example, if the target object's color was located at 0° on the color wheel and the competing object's color was at 24°, a color response at 350° would be considered away from the competing object's color (see Fig. 1c and the Method section). We defined a repulsion effect as occurring for a condition if the mean percentage of "away" responses was greater than 50% (i.e., that most color reports were biased away from the color of the competing object).

For the Day 1 posttraining test, an analysis of variance (ANOVA) with similarity condition as a factor revealed a robust main effect of similarity on the percentage of responses away from the competitor, F(2,44) = 10.11, p < .001, $\eta_G^2 = .22$ (Fig. 2c). Critically, there was a strong repulsion effect in the high-similarity condition (*M* = 60.80%, *SD* = 12.00%), *t*(22) = 4.31, *p* < .001, 95% confidence interval (CI) for the difference from 50% = [5.61, 15.99], Cohen's d = 1.27, but not in the moderate-similarity condition (M = 48.70%, SD = 9.37%), t(22) = -0.67, p = .512, 95% CI for the difference from 50% = [-5.36, 2.75], Cohen's d = -0.20, or low-similarity condition (M = 49.93%, SD = 10.16%), t(22) = -0.03,p = .973, 95% CI for the difference from 50% = [-4.46,4.32], Cohen's d = -0.01. Thus, when similarity between competing objects was high, color memory for a target object was systematically biased away from the color of the competing object. Follow-up comparisons confirmed that the percentage of "away" responses was significantly greater in the high-similarity condition compared with both the moderate-similarity condition, t(22) = 4.47, p < .001, 95% CI for the difference between conditions = [6.49, 17.71], Cohen's d = 1.12, and lowsimilarity condition, t(22) = 3.80, p = .001, 95% CI for the difference between conditions = [4.93, 16.81], Cohen's d = 0.98. The selectivity of the repulsion effect to the high-similarity condition is striking when one considers that interference errors during the associative-memory tests in the training rounds were much more common in the high-similarity condition than in the moderate- or low-similarity conditions (see Fig. S2).

Interestingly, the repulsion effect strongly persisted the following day: The main effect of similarity condition was again significant at the Day 2 posttraining test, F(2, 42) = 9.82, p < .001, $\eta_G^2 = .19$ (Fig. 2c), and there was a repulsion effect in the high-similarity condition (M = 59.77%, SD = 11.17%), t(21) = 4.11, p < .001, 95% CI for the difference from 50% = [4.82, 14.72], Cohen's d = 1.24, but not in the moderate-similarity condition (M = 47.80%, SD = 10.55%), t(21) = -0.98, p = .340, 95% CI for the difference from 50% = [-6.87, 2.48], Cohen's d = -0.29, or low-similarity condition (M = 49.39%,SD = 12.05%), t(21) = -0.24, p = .816, 95% CI for the difference from 50% = [-5.95, 4.73], Cohen's d = -0.07.

Experiment 2

In a replication study (Experiment 2), we used the same procedure as in Experiment 1 except for a few minor changes (see the Method section for details and rationale). In particular, we reduced the number of color-memory tests during the training rounds by 50% and eliminated the Day 2 posttraining test. Performance across the training rounds is reported in Figures 2a and 2b and in Figures S1 and S2. As in Experiment 1, across training rounds, mean color-memory error decreased, F(1, 34) = 516.4, p < .001, $\eta_G^2 = .83$, and accuracy on the associative-memory tests increased, F(1, 35) = 690.4, p < .001, $\eta_G^2 = .85$. Accuracy differed across color-similarity condition, F(2, 70) = 18.77, p < .001, $\eta_G^2 = .19$.

Also as in Experiment 1, the percentage of responses away from the competitor during the color-memory posttraining test robustly varied across color-similarity conditions, F(2, 70) = 6.79, p = .002, $\eta_G^2 = .09$ (Fig. 2c). Critically, we again observed a significant repulsion effect in the high-similarity condition (M = 54.63%, SD = 12.74%, t(35) = 2.18, p = .036, 95% CI for the difference from 50% = [0.32, 8.95], Cohen's d = 0.51. The percentage of "away" responses did not significantly differ from 50% in the moderate-similarity condition (M = 47.95%, SD = 10.42%), t(35) = -1.18, p = .247, 95% CI for the difference from 50% = [-5.57, 1.48], Cohen's d = -0.28, and there was a trend in the opposite direction (below 50%) in the low-similarity condition (M = 46.05%, SD = 12.07%), t(35) = 1.96, p = .057, 95%CI for the difference from 50% = [-8.04, 0.13], Cohen's d = -0.46. Follow-up tests confirmed that the percentage of "away" responses was again significantly higher in the high-similarity condition compared with both the moderate-similarity condition, t(35) = 3.10, p = .004, 95% CI for the difference between conditions = [2.30, 11.06], Cohen's d = 0.57, and the low-similarity condition, t(35) =3.36, p = .002, 95% CI for the difference between conditions = [3.40, 13.77], Cohen's d = 0.69. Thus, as in Experiment 1, we observed a selective repulsion effect in color memory specifically when there was high similarity between competing objects.

Experiment 3

Experiments 1 and 2 strongly established that the repulsion effect is competition dependent, as it was selective to the high-similarity condition. Interestingly, however, although hippocampal repulsion effects are also competition dependent, the relationship between competition and repulsion is thought to be nonmonotonic: that is, with sufficiently strong competition, representations will fail to diverge (Ritvo et al., 2019). In Experiment 3, we



Fig. 3. Ultra-similarity condition and results from all conditions of Experiment 3. Experiment 3 dropped the high-similarity condition used in the previous experiments and added an ultra-similarity condition (a; 6° apart). An example of competing images from the ultra-similarity condition is shown. Mean color-memory error (b; absolute distance between reported and target color values from the color-memory tests) and accuracy on the associative-memory test (c; percentage of trials for which the target face was selected; chance = 25%) are shown as a function of training round and similarity condition. The percentage of responses away from the competitor on the posttraining test (d) is shown as a function of training day and similarity condition. Large dots represent means, and small dots represent data from individual subjects. The asterisk represents a significant difference (p < .05) from 50%. Error bars represent standard errors of the mean.

tested whether there was a nonmonotonic relationship between color similarity and repulsion by shifting the range of color similarity to include a moderate-similarity condition (again, 48°), a high-similarity condition (again, 24°) and a new ultra-similarity condition (6°; Fig. 3a).

Performance across the training rounds is reported in Figures 3b and 3c and in Figures S1 and S2. Across training rounds, mean color-memory error decreased, $F(1, 37) = 186.5, p < .001, \eta_G^2 = .70$, and accuracy on the associative-memory tests increased, F(1, 37) = 326.9, p < .001, $\eta_G^2 = .75$. Accuracy differed across colorsimilarity conditions, F(2, 74) = 129.9, p < .001, $\eta_{G}^{2} =$.60. Of particular relevance, accuracy on the associativememory test (during the training rounds) was significantly lower in the ultra-similarity condition than in the high-similarity condition, t(37) = -11.39, p < .001, 95% CI for the difference between conditions = [-17.95], -12.53], Cohen's d = -1.85, and in the moderatesimilarity condition, t(37) = -16.26, p < .001, 95% CI for the difference between conditions = [-24.26, -18.86], Cohen's d = -2.96 (Fig. 3c; see also Fig. S2). Nonetheless, in the last round of the associative-memory test, subjects selected the target faces at above-chance rates in all similarity conditions (chance = 25%, all means > 66%, ps < .001). Thus, the ultra-similarity condition clearly increased interference relative to the highsimilarity condition, but subjects were still generally successful at memory-based discrimination between these extremely similar colors.

Results from the posttraining test again revealed that color similarity influenced the percentage of responses away from the competitor, F(2, 74) = 5.45, p = .006, $\eta_G^2 = .06$

(Fig. 3d). However, the relationship between similarity and repulsion followed the predicted nonmonotonic pattern. As in Experiments 1 and 2, there was a significant repulsion effect in the high-similarity condition (M = 54.74%, SD = 13.42%), t(37) = 2.17, p = .036, 95%CI for the difference from 50% = [0.32, 9.15], Cohen's d = 0.50; no repulsion effect in the moderate-similarity condition (M = 47.19%, SD = 10.93%), t(37) = -1.58, p =.122, 95% CI for the difference from 50% = [-6.40, 0.79], Cohen's d = -0.36; and a significant difference between the high- and moderate-similarity conditions, t(37) =3.04, p = .004, 95% CI for the difference between conditions = [2.51, 12.58], Cohen's d = 0.61. In the ultrasimilarity condition, however, the percentage of "away" responses did not differ from 50% (M = 50.75%, SD =(8.52%), t(37) = 0.54, p = .593, 95% CI for the difference from 50% = [-2.06, 3.55], Cohen's d = 0.12, confirming that, with sufficiently high similarity, the repulsion effect was eliminated. Although the percentage of "away" responses was numerically lower in the ultrasimilarity condition than in the high-similarity condition, this difference was not statistically significant, t(37) = -1.63, p = .111, 95% CI for the difference between conditions = [-0.10, 8.94], Cohen's d = 0.35. Interestingly, despite the much higher rate of interference errors in the ultra-similarity condition compared with the moderate-similarity condition (see Fig. S2), the percentage of color responses away from the competing object's color was numerically higher in the ultrasimilarity condition than in the moderate-similarity condition, but this difference was not statistically significant, t(37) = 1.89, p = .067, 95% CI for the difference between

conditions = [-0.026, 7.37], Cohen's d = 0.36. Taken together, performance across the three similarity conditions suggests a local maximum in the repulsion effect that occurred when similarity was high (24°) but not too high (6°).

Relationship between repulsion and memory interference

Thus far, we have shown that the repulsion effect is triggered by similarity among memories. This raises the complementary question: What is the consequence of repulsion? From an adaptive perspective, repulsion may carry the important benefit that by exaggerating the differences between similar memories, it serves to reduce memory interference. To test for a relationship between repulsion and interference, we considered data from Experiments 1 through 3 and focused specifically on the high-similarity condition because a significant repulsion effect was observed in this condition across all three experiments (for analyses of moderateand low-similarity conditions, see Figs. S4 and S5 in the Supplemental Material). For each subject in each experiment, we computed (a) the mean percentage of responses away from the competitor on the basis of data from the immediate posttraining test and (b) the mean number of interference errors across the last three rounds of the associative-memory test (during the training rounds). As a first step, we tested for across-subjects correlations between the mean percentage of "away" responses on the color-memory posttraining test and mean interference errors on the associative-memory test. Strikingly, a significant negative correlation was observed for each experiment—Experiment 1: r(21) =-.61, p = .002; Experiment 2: r(34) = -.51, p = .001; Experiment 3: r(36) = -.44, p = .006 (see Figs. 4a and 4b). Thus, stronger color-memory repulsion was associated with fewer interference-related errors during the associative-memory test. We observed a similar negative relationship between repulsion and interference errors in a within-subjects, item-level analysis (see Fig. S6 in the Supplemental Material).

One potential caveat with the correlations described above is that they may partly reflect that subjects who suffered more interference errors during associativememory tests were more likely to remember the wrong color (i.e., the competing object's color) during the color-memory posttraining test. From this perspective, it is possible that all subjects showed comparable levels of repulsion when they recalled the correct color, but subjects with more interference errors also recalled the wrong color with greater frequency, thereby pulling down their percentage of responses away from the competitor. To address this concern, we performed a second, more targeted analysis that focused on the distribution of correct color-memory responses. Using a median split, we first divided all subjects (within each experiment) into high-interference and low-interference groups on the basis of the mean number of interference errors during the last three associative-memory tests in the training rounds (high-interference group: M = 18.45%, SD = 7.58%; low-interference group: M = 2.12%, SD =2.59%). We then computed the frequency of responses on the posttraining test that fell in each of four colorspace bins. Two of these bins were centered around the target color value (±11° from the target color), and two of these bins were centered around the competitor color value $(\pm 11^{\circ}$ from the competitor color; Fig. 4c). This allowed us to separate color-memory responses that were correct (±11° from the target) from swap errors (±11° from the competitor). Note that responses that were precisely equal to the target or competitor color were excluded from this analysis.

Not surprisingly, high-interference subjects tended to commit more swap errors in the color-memory test (M = 21.47%, SD = 9.51%) than did low-interference subjects (M = 10.07%, SD = 7.28%), t(82) = 6.17, p < 6.17.001, 95% CI for the difference between conditions = [7.73, 15.10], Cohen's d = 1.35. Of critical interest, however, was the distribution of correct responses-that is, whether there was a difference in the frequency of correct "toward" and correct "away" responses. For highinterference subjects, correct color-memory responses were evenly distributed around the actual target value (no difference in frequency of correct-toward vs. correct-"away" responses), F(1, 39) = 2.07, p = .158, $\eta_{G}^{2} = .03$ (see the left side of Fig. 4c). For low-interference subjects, however, there was a strong asymmetry, with significantly more responses in the correct-"away" bin compared with the correct-"toward" bin, *F*(1, 39) = 58.00, *p* < .001, $\eta_G^2 = .39$ (see the right side of Fig. 4c). The interaction between subjects (high interference vs. low interference) and response bins (correct "toward" vs. correct "away") was also significant, F(1, 78) = 13.03, p < .001, $\eta_G^2 = .08$. None of these effects interacted with experiment number (ps > .4). Thus, even when restricting focus to color-memory responses that were correct (i.e., removing swap errors), we found clear evidence for an adaptive distortion of color memory: Subjects who made the fewest interference errors during the associativememory test exhibited a robust repulsion effect whereby color memory was systematically biased away from the color of the competing object.

Experiment 4

In Experiments 1 through 3, the associative-memory task during the training rounds explicitly required subjects to discriminate between the competing objects. In Experiment 4, we tested whether this discrimination



Fig. 4. Analysis of data from Experiments 1 through 3. The scatterplot in (a) shows across-subject correlations between mean percentage of posttraining test (Day 1) color-memory responses that were away from the competitor and mean percentage of interference errors during the last three rounds of the associative-memory test (high-similarity condition only). The scatterplot in (b) shows across-subject correlations for the same variables, except that each measure was *z*-scored within experiment, allowing for a single correlation to be calculated for the pooled data. Interference errors were defined as selecting the face associated with the competitor object. Lines show best-fitting regressions, and shaded areas around the lines indicate 68% confidence intervals. The graphs in (c) show the distribution of correct and swap responses both toward and away from the competitor, separately for participants who showed a high number and a low number of interference errors during the last three rounds of the associative-memory tests. Asterisks indicate a significant difference (p < .001) between "correct" responses. Error bars represent standard errors of the mean.



Fig. 5. Inference test and results from Experiment 4. All procedures of Experiment 4 were identical to those of Experiment 1 except for a subtle change to the associative-memory test during the training rounds. Instead of requiring subjects to discriminate between competing colors (as in Experiment 1), the associative-memory test consisted of an inference test that required subjects to generalize across competing objects. On each trial of the inference test, a probe face was presented, and subjects had to select which face, from a set of four options, was associated with the same object as the probe face (irrespective of color). Thus, what was previously the competitor face in Experiments 1, 2 and 3 was now the correct response. Mean color-memory error (b; absolute distance between reported and target color values from the color-memory tests) and accuracy on the inference test (c; percentage of trials for which the target face was selected; chance = 25%) are shown as a function of training round and similarity condition. The percentage of responses away from the competitor on the posttraining test (d) is shown as a function of training day and similarity condition. Large dots represent means, and small dots represent data from individual subjects. Asterisks represent significant differences (*p < .05, **p < .001) from 50%. Error bars represent standard errors of the mean.

demand was necessary for inducing repulsion. The critical difference in Experiment 4, relative to Experiments 1 through 3, was that we changed the procedures for the associative-memory test in the training rounds so that it now promoted integration across overlapping associations (Richter et al., 2016; Shohamy & Wagner, 2008; Zeithamova et al., 2012). Specifically, the new associative-memory test was an inference test (Zeithamova et al., 2012) that required subjects to generalize across overlapping associations. On each trial in the inference test, a face image (probe) was presented at the top of the screen, and subjects had to select a matching face from a set of four options presented below (Fig. 5a). A matching face was defined as a face associated with the same object category as the probe (ignoring differences in color). For example, two faces would match if they were each associated with a backpack, despite the fact that the backpacks differed in color. Thus, although the inference test still required associative learning (i.e., object-face learning), it did not require discriminating between similar objects. However, because color memory was still tested during the training rounds (as in all prior experiments), color memory remained relevant, and subjects showed robust improvements in color memory across training rounds, F(1, 25) =101.03, p < .001, $\eta_G^2 = .63$ (Fig. 5b).

Subjects performed well in the inference task, with a significant improvement in performance across training rounds, F(1, 25) = 225.3, p < .001, $\eta_{G}^2 = .79$ (Fig. 5c).

However, in contrast to what we found in the associativememory tests in Experiments 1 through 3, performance in the inference test no longer differed as a function of color similarity, F(2, 50) = 0.42, p = .659, $\eta_G^2 = .006$ (Fig. 5c). That is, the ability to generalize across associations with common object categories was not influenced by color similarity.

As in all of the prior experiments, the percentage of responses away from the competitor during the colormemory posttraining test varied by similarity condition, $F(2, 50) = 5.21, p = .009, \eta_G^2 = .09$ (Fig. 5d). However, for the high-similarity condition the percentage of "away" responses no longer differed from 50%, t(25) =-0.49, p = .626, 95% CI for the difference from 50% =[-6.64, 4.07], Cohen's d = -0.14. Interestingly, the percentage of "away" responses was significantly lower than 50% in both the moderate- and low-similarity conditions—moderate: t(25) = -2.08, p = .048, 95% CI for the difference from 50% = [-9.83, -0.04], Cohen's d =-0.58; low: t(25) = -5.17, p < .001, 95% CI for the difference from 50% = [-13.99, -6.02], Cohen's d = -1.43 suggesting an attraction effect. A direct comparison of the percentage of "away" responses in the highsimilarity conditions in Experiment 4 and Experiment 1-the experiment most closely matched to Experiment 4—revealed a significant difference, t(47) = 3.33, p =.002, 95% CI for the difference between conditions = [4.77, 19.39], Cohen's d = 0.95, confirming that the subtle change in task demands (encouraging integration as opposed to discrimination) significantly reduced the repulsion effect.

Discussion

Although numerous studies have documented the situations and contexts in which similarity between events leads to interference and forgetting (Anderson, 2003; Barnes & Underwood, 1959; Mensink & Raaijmakers, 1988; Osgood, 1949), here we considered whether similarity between events triggers adaptive distortions in how those events are remembered. We have shown that interference between similar-colored objects induces a repulsion effect in which the colors of these objects are remembered as being further apart than they actually are. This repulsion effect was highly dependent on competition (color similarity) and was sensitive to task demands. Critically, repulsion was also adaptive: Greater repulsion was strongly associated with fewer interference-related errors during associative-memory retrieval. These findings provide striking evidence of adaptive-memory distortions that are triggered by competition among highly similar memories.

Our core predictions were motivated by-and our findings strongly parallel-recent evidence of repulsion effects in human hippocampal-activity patterns (Ballard et al., 2019; Chanales et al., 2017; Favila et al., 2016; Hulbert & Norman, 2015; Schapiro et al., 2012). For example, Chanales et al. found that hippocampal repulsion was greatest for the segments of spatial routes that were most difficult to discriminate, suggesting that similarity triggered hippocampal repulsion. Likewise, we found robust and selective evidence of behavioral-memory repulsion (Experiments 1–3) when color similarity—and memory interference-was high (24° apart). Notably, this finding is precisely opposite to what might be expected on the basis of classic interference studies. Specifically, from an interference perspective, subjects should have been more likely to confuse the colors of overlapping objects when color similarity was high. This would translate to a color-memory bias toward competing objects, or an *attraction effect*. Although we did observe attraction effects during early stages of learning (see Fig. S1)-entirely consistent with interferencethese attraction effects switched to repulsion with extended training, which is exactly what has been observed with hippocampal repulsion (Chanales et al., 2017; Favila et al., 2016; Schlichting et al., 2015). Consequently, repulsion can be viewed as an aftereffect of interference.

Our findings also parallel prior evidence that hippocampal repulsion plays an adaptive role in reducing interference (confusability) among highly similar memories (Colgin et al., 2008; Favila et al., 2016; Hulbert & Norman, 2015). Specifically, we have shown that color-memory repulsion was overwhelmingly more pronounced in subjects that suffered the fewest interference errors during associative-memory retrieval. When specifically considering high-similarity trials with correct color memory (defined as $\pm 11^{\circ}$ of the target), we found a striking difference between subjects with high and low rates of interference errors: Subjects who had more interference errors had response distributions that were centered on the veridical color value; in contrast, subjects with fewer interference errors exhibited a response distribution that was shifted away from the color of the competing object. Thus, lower associative interference was associated with a systematic adaptive distortion in how memory cues were remembered.

In order to induce a repulsion effect in color memory, we deliberately developed a training procedure that involved alternation between study and competitive retrieval (Experiments 1-3). This procedure was inspired by evidence that study/retrieval alternation is very effective in creating distinct representations of overlapping memories (Hulbert & Norman, 2015; Storm et al., 2008) and in inducing differentiation of hippocampal-activity patterns (Hulbert & Norman, 2015; Kim et al., 2017). These dynamics have also been detailed in computational models that address how episodic-memory interference is resolved (Norman et al., 2006, 2007). More generally, our results build on evidence that competitive remembering triggers active mechanisms that reshape the memory landscape in order to reduce interference (Anderson, 2003; Anderson et al., 1994; Levy & Anderson, 2002; Norman et al., 2006, 2007).

Across our experiments, we identified several boundary conditions for the repulsion effect. First, as noted above, we consistently observed repulsion in the highsimilarity condition (24°) but not in the moderate- or low-similarity conditions (48° or 72°). However, with even higher similarity (6°), the repulsion effect was no longer significant. Thus, the relationship between similarity and the repulsion effect followed an inverted-Ushape function, suggesting a sweet spot at which repulsion occurs. This finding is consistent with theoretical perspectives on the relationship between memory competition and memory plasticity (Ritvo et al., 2019). Specifically, memory representations are thought to be most susceptible to plasticity (weakening or distortion) at particular levels of competition. If memory representations are too similar or too dissimilar, then plasticity is not expected to occur. This theoretical perspective is supported by several examples of nonmonotonic relationships between neural measures of competition and memory or plasticity (Chanales et al., 2017; Detre et al., 2013; Lewis-Peacock & Norman, 2014; Newman & Norman, 2010).

Another boundary condition to the repulsion effect relates to task demands. The repulsion effect was not observed when task demands explicitly encouraged integration (instead of discrimination) of similar objects (Experiment 4). Interestingly, this integration demand led to an attraction effect for the low- and moderatesimilarity conditions but not for the high-similarity condition. On the one hand, storing a single averaged color value for each object pair (i.e., attraction) would seemingly be an efficient strategy when task demands require integration (Gluck & Myers, 1993; Richards et al., 2014). However, it is possible-though this is speculative-that event similarity triggers some degree of repulsion regardless of task demands (Favila et al., 2016) and that, in Experiment 4, results in the highsimilarity condition reflect offsetting effects of integration and repulsion. Detailed consideration of this point is beyond the scope of the present article, but our findings establish that task demands are an important factor, along with event similarity. Additional research will be required to map out exactly how and when repulsion effects are influenced by task demands.

Although our findings were specifically motivated by empirical phenomena and theoretical perspectives in the field of episodic memory, they contribute to a broader literature documenting adaptive exaggeration in cognitive processes. For example, similar biases have previously been documented in visual working memory (Bae & Luck, 2017; Rademaker et al., 2015), estimates of temporal duration (Ezzyat & Davachi, 2014), and judgments of social categories (Förster et al., 2008; Krueger & Rothbart, 1990). This raises the question of whether the repulsion effect we observed is, fundamentally, a bias in episodic memory or whether the bias might occur during another cognitive-processing stage. In particular, it is possible that the bias occurred during perception and was then reinstated during memory retrieval. This framing is not incompatible with our claims. That said, it is important to emphasize that any bias during perception would still be dependent on long-term memory: A perceptual bias could occur only to the extent that a remembered stimulus exerted an influence on a currently perceived stimulus (Teng & Kravitz, 2019). Moreover, it is interesting to note that damage to the hippocampus (a structure critical for episodic-memory formation) is also associated with impairments in fine-grained perceptual discriminations (Aly et al., 2013), suggesting that the distinction between memory and perception may not be categorical (Aly & Turk-Browne, 2018). Ultimately, although it is an interesting question whether the repulsion effect reported here also occurred during perception, the critical points are that the repulsion effect we report (a) was induced by long-term memory, (b) was remarkably stable over time (e.g., it persisted for ~24 hr in Experiment 1), and (c) strongly predicted associative-interference errors in a canonical episodic-memory paradigm.

Collectively, our results robustly establish that similarity between long-term memories triggers a repulsion in remembered feature values and that this exaggeration of remembered features is highly adaptive. These findings strongly support the idea that memory distortions generally reflect the operation of an adaptivememory system (Schacter, 1999) while providing new evidence of how such distortions can specifically mitigate memory interference.

Transparency

Action Editor: D. Stephen Lindsay Editor: D. Stephen Lindsay

Author Contributions

A. J. H. Chanales and B. A. Kuhl designed the experiments. A. G. Tremblay-McGaw supervised or conducted all data collection. A. J. H. Chanales, M. L. Drascher, and B. A. Kuhl analyzed the data. All authors contributed to the writing and editing of the manuscript and approved the final version for submission.

Declaration of Conflicting Interests

The author(s) declared that there were no conflicts of interest with respect to the authorship or the publication of this article.

Funding

This work was supported by Grant No. R01-NS089729 from the National Institute of Neurological Disorders and Stroke to B. A. Kuhl.

Open Practices

Stimuli, data, and relevant code for Experiments 1 through 4 have been made publicly available via OSF and can be accessed at https://osf.io/d4pmg/. The design and analysis plans for the experiments were not preregistered. This article has received the badges for Open Data and Open Materials. More information about the Open Practices badges can be found at http://www.psychologicalscience .org/publications/badges.



ORCID iD

Brice A. Kuhl i https://orcid.org/0000-0001-5229-5400

Supplemental Material

Additional supporting information can be found at http://journals.sagepub.com/doi/suppl/10.1177/0956797620972490

References

Aly, M., Ranganath, C., & Yonelinas, A. P. (2013). Detecting changes in scenes: The hippocampus is critical for strength-based perception. *Neuron*, 78, 1127–1137. https:// doi.org/10.1016/j.neuron.2013.04.018

- Aly, M., & Turk-Browne, N. B. (2018). Flexible weighting of diverse inputs makes hippocampal function malleable. *Neuroscience Letters*, 680, 13–22. https://doi.org/ 10.1016/j.neulet.2017.05.063
- Anderson, M. C. (2003). Rethinking interference theory: Executive control and the mechanisms of forgetting. *Journal of Memory and Language*, 49, 415–445. https:// doi.org/10.1016/j.jml.2003.08.006
- Anderson, M. C., Bjork, R. A., & Bjork, E. L. (1994). Remembering can cause forgetting: Retrieval dynamics in long-term memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20*, 1063–1087. https:// doi.org/10.1037/0278-7393.20.5.1063
- Bae, G.-Y., & Luck, S. J. (2017). Interactions between visual working memory representations. *Attention, Perception,* & *Psychophysics*, 79, 2376–2395. https://doi.org/10.3758/ s13414-017-1404-8
- Bakker, A., Kirwan, C. B., Miller, M., & Stark, C. E. L. (2008). Pattern separation in the human hippocampal CA3 and dentate gyrus. *Science*, *319*, 1640–1642. https://doi .org/10.1126/science.1152882
- Ballard, I. C., Wagner, A. D., & McClure, S. M. (2019). Hippocampal pattern separation supports reinforcement learning. *Nature Communications*, 10, Article 1073. https://doi.org/10.1038/s41467-019-08998-1
- Barnes, J. M., & Underwood, B. J. (1959). "Fate" of first-list associations in transfer theory. *Journal of Experimental Psychology*, 58, 97–105. https://doi.org/10.1037/h0047507
- 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. https://doi.org/10.1037/a0029649
- Chanales, A. J. H., Oza, A., Favila, S. E., & Kuhl, B. A. (2017). Overlap among spatial memories triggers repulsion of hippocampal representations. *Current Biology*, 27, 2307– 2317.e5. https://doi.org/10.1016/j.cub.2017.06.057
- Colgin, L. L., Moser, E. I., & Moser, M.-B. (2008). Understanding memory through hippocampal remapping. *Trends in Neurosciences*, 31, 469–477. https://doi.org/10.1016/j .tins.2008.06.008
- Detre, G. J., Natarajan, A., Gershman, S. J., & Norman, K. A. (2013). Moderate levels of activation lead to forgetting in the think/no-think paradigm. *Neuropsychologia*, 51, 2371–2388. https://doi.org/10.1016/j.neuropsycholo gia.2013.02.017
- Dimsdale-Zucker, H. R., Ritchey, M., Ekstrom, A. D., Yonelinas, A. P., & Ranganath, C. (2018). CA1 and CA3 differentially support spontaneous retrieval of episodic contexts within human hippocampal subfields. *Nature Communications*, *9*, Article 294. https://doi.org/10.1038/ s41467-017-02752-1
- Ezzyat, Y., & Davachi, L. (2014). Similarity breeds proximity: Pattern similarity within and across contexts is related to later mnemonic judgments of temporal proximity. *Neuron*, 81, 1179–1189. https://doi.org/10.1016/j.neu ron.2014.01.042
- Favila, S. E., Chanales, A. J. H., & Kuhl, B. A. (2016). Experience-dependent hippocampal pattern differentiation

prevents interference during subsequent learning. *Nature Communications*, 7, Article 11066. https://doi .org/10.1038/ncomms11066

- Förster, J., Liberman, N., & Kuschel, S. (2008). The effect of global versus local processing styles on assimilation versus contrast in social judgment. *Journal of Personality and Social Psychology*, 94, 579–599. https://doi .org/10.1037/0022-3514.94.4.579
- Gluck, M. A., & Myers, C. E. (1993). Hippocampal mediation of stimulus representation: A computational theory. *Hippocampus*, *3*, 491–516. https://doi.org/10.1002/hipo .450030410
- Hulbert, J. C., & Norman, K. A. (2015). Neural differentiation tracks improved recall of competing memories following interleaved study and retrieval practice. *Cerebral Cortex*, 25, 3994–4008. https://doi.org/10.1093/cercor/bhu284
- Kim, G., Norman, K. A., & Turk-Browne, N. B. (2017). Neural differentiation of incorrectly predicted memories. *The Journal of Neuroscience*, 37, 2022–2031. https://doi .org/10.1523/JNEUROSCI.3272-16.2017
- Krueger, J., & Rothbart, M. (1990). Contrast and accentuation effects in category learning. *Journal of Personality* and Social Psychology, 59, 651–663. https://doi.org/ 10.1037/0022-3514.59.4.651
- Levy, B., & Anderson, M. C. (2002). Inhibitory processes and the control of memory retrieval. *Trends in Cognitive Sciences*, 6, 299–305. https://doi.org/10.1016/S1364-6613(02)01923-X
- Lewis-Peacock, J. A., & Norman, K. A. (2014). Competition between items in working memory leads to forgetting. *Nature Communications*, 5, Article 5768. https://doi.org/ 10.1038/ncomms6768
- Mensink, G.-J., & Raaijmakers, J. G. (1988). A model for interference and forgetting. *Psychological Review*, 95, 434–455. https://doi.org/10.1037/0033-295X.95.4.434
- Newman, E. L., & Norman, K. A. (2010). Moderate excitation leads to weakening of perceptual representations. *Cerebral Cortex*, 20, 2760–2770. https://doi.org/10.1093/ cercor/bhq021
- Norman, K. A., Newman, E. L., & Detre, G. (2007). A neural network model of retrieval-induced forgetting. *Psychological Review*, 114, 887–953. https://doi.org/10.1037/0033-295X.114.4.887
- Norman, K. A., Newman, E. L., Detre, G., & Polyn, S. (2006). How inhibitory oscillations can train neural networks and punish competitors. *Neural Computation*, 18, 1577–1610. https://doi.org/10.1162/neco.2006.18.7.1577
- O'Reilly, R. C. (2004). The division of labor between the neocortex and hippocampus. In G. Houghton (Ed.), *Connectionist models in cognitive psychology* (pp. 143–174). Psychology Press.
- Osgood, C. E. (1949). The similarity paradox in human learning: A resolution. *Psychological Review*, *56*(3), 132–143. https://doi.org/10.1037/h0057488
- 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. https://doi.org/10.1037/xhp0000110

- Richards, B. A., Xia, F., Santoro, A., Husse, J., Woodin, M. A., Josselyn, S. A., & Frankland, P. W. (2014). Patterns across multiple memories are identified over time. *Nature Neuroscience*, 17, 981–986. https://doi.org/10.1038/nn.3736
- Richter, F. R., Chanales, A. J. H., & Kuhl, B. A. (2016). Predicting the integration of overlapping memories by decoding mnemonic processing states during learning. *NeuroImage*, 124, 323–335. https://doi.org/10.1016/j.neu roimage.2015.08.051
- Ritvo, V. J. H., Turk-Browne, N. B., & Norman, K. A. (2019). Nonmonotonic plasticity: How memory retrieval drives learning. *Trends in Cognitive Sciences*, 23, 726–742. https://doi.org/10.1016/j.tics.2019.06.007
- Schacter, D. L. (1999). The seven sins of memory: Insights from psychology and cognitive neuroscience. *American Psychologist*, 54, 182–203. https://doi.org/10.1037/0003-066X.54.3.182
- Schapiro, A. C., Kustner, L. V., & Turk-Browne, N. B. (2012). Shaping of object representations in the human medial temporal lobe based on temporal regularities. *Current Biology*, 22, 1622–1627. https://doi.org/10.1016/j.cub.2012.06.056
- Schlichting, M. L., Mumford, J. A., & Preston, A. R. (2015). Learning-related representational changes reveal

dissociable integration and separation signatures in the hippocampus and prefrontal cortex. *Nature Communications*, *6*, Article 8151. https://doi.org/10.1038/ncomms9151

- Shohamy, D., & Wagner, A. D. (2008). Integrating memories in the human brain: Hippocampal-midbrain encoding of overlapping events. *Neuron*, 60, 378–389. https://doi .org/10.1016/j.neuron.2008.09.023
- Storm, B. C., Bjork, E. L., & Bjork, R. A. (2008). Accelerated relearning after retrieval-induced forgetting: The benefit of being forgotten. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34, 230–236. https:// doi.org/10.1037/0278-7393.34.1.230
- Teng, C., & Kravitz, D. J. (2019). Visual working memory directly alters perception. *Nature Human Behaviour*, *3*, 827–836. https://doi.org/10.1038/s41562-019-0640-4
- Yassa, M. A., & Stark, C. E. L. (2011). Pattern separation in the hippocampus. *Trends in Neurosciences*, 34, 515–525. https://doi.org/10.1016/j.tins.2011.06.006
- Zeithamova, D., Dominick, A. L., & Preston, A. R. (2012). Hippocampal and ventral medial prefrontal activation during retrieval-mediated learning supports novel inference. *Neuron*, 75, 168–179. https://doi.org/10.1016/j .neuron.2012.05.010