

When past is present: Substitutions of long-term memory for sensory evidence in perceptual judgments

Judith E. Fan*

Department of Psychology, Princeton University,
Princeton, NJ, USA



J. Benjamin Hutchinson*

Department of Psychology, Princeton University,
Princeton, NJ, USA
Princeton Neuroscience Institute, Princeton University,
Princeton, NJ, USA



Nicholas B. Turk-Browne

Department of Psychology, Princeton University,
Princeton, NJ, USA
Princeton Neuroscience Institute, Princeton University,
Princeton, NJ, USA



When perception is underdetermined by current sensory inputs, memories for related experiences in the past might fill in missing detail. To evaluate this possibility, we measured the likelihood of relying on long-term memory versus sensory evidence when judging the appearance of an object near the threshold of awareness. Specifically, we associated colors with shapes in long-term memory and then presented the shapes again later in unrelated colors and had observers judge the appearance of the new colors. We found that responses were well characterized as a bimodal mixture of original and current-color representations (vs. an integrated unimodal representation). That is, although irrelevant to judgments of the current color, observers occasionally anchored their responses on the original colors in memory. Moreover, the likelihood of such memory substitutions increased when sensory input was degraded. In fact, they occurred even in the absence of sensory input when observers falsely reported having seen something. Thus, although perceptual judgments intuitively seem to reflect the current state of the environment, they can also unknowingly be dictated by past experiences.

expectations about what is likely to appear in a given context and how it will look. The supplemental role of memory in perception might be especially needed under impoverished viewing conditions (e.g., dim lighting; von Helmholtz, 1867). For instance, in a dark movie theater, it may be easier to recognize the face of a close friend than that of an acquaintance, even though it may not be difficult to identify either under normal illumination.

There are multiple ways in which past experience can influence present perceptual performance. Benefits may derive from extensive practice on specific perceptual tasks, as in the case of perceptual learning (Sagi, 2011). For example, consistent exposure to dots moving in one direction leads to enhanced motion perception in that direction (Ball & Sekuler, 1987; Watanabe, Náñez, & Sasaki, 2001). Perceptual expectations formed on the basis of this exposure can bias judgments of stimulus features (e.g., motion direction) toward the average of feature values experienced so far (Huang & Sekuler, 2010), or those most frequently experienced, and even on trials on which no stimulus is present (Chalk, Seitz, & Series, 2010). Moreover, over a lifetime of experience, certain objects (e.g., bananas) are perceived to have their canonical color even when presented achromatically (Hansen, Olkkonen, Walter, & Gegenfurtner, 2006).

Benefits may also accrue from minimal experience, such as a single exposure to a word or object, which results in faster and more accurate identification upon repetition (Tulving & Schacter, 1990), even lasting for

Introduction

What we see is intuitively thought to be determined by the light reaching our eyes, but it is also influenced by memories for what we have seen before. Indeed, past experience provides a rich source of knowledge and

Citation: Fan, J. E., Hutchinson, J. B., & Turk-Browne, N. B. (2016). When past is present: Substitutions of long-term memory for sensory evidence in perceptual judgments. *Journal of Vision*, 16(8):1, 1–12, doi:10.1167/16.8.1.

doi: 10.1167/16.8.1

Received June 19, 2015; published June 1, 2016

ISSN 1534-7362



weeks (Cave, 1997). One remarkable feature of such priming is that it does not require awareness that an item has been seen before. When encountering a repeated stimulus, people may erroneously attribute their perceptual fluency to aspects of current viewing conditions (e.g., presentation duration, the level of noise) rather than to memory (Witherspoon & Allan, 1985; Jacoby, Allan, & Collins, 1988). Furthermore, amnesic patients, who are not able to consciously recall recent experiences, nevertheless can exhibit priming (Warrington & Weiskrantz, 1974).

Taken together, this prior work shows that information encoded into long-term memory can automatically influence perception, partly by enhancing perceptual sensitivity for that information when it is presented again. Although this suggests that long-term memory interacts with current sensory inputs in some way to determine perception, *how* these interactions occur is less well understood.

Recently, the field of visual memory has developed quantitative modeling techniques to characterize the strength and fidelity of visual representations (Zhang & Luck, 2008; van den Berg, Shin, Chou, George, & Ma, 2012; Fougnie, Suchow, & Alvarez, 2012; Suchow, Brady, Fougnie, & Alvarez, 2013). These tools have spurred advances in our understanding of diverse phenomena (see Brady, Konkle, & Alvarez, 2011), including feature binding in visual working memory (e.g., Fougnie, Asplund, & Marois, 2010; Bays, Wu, & Husain, 2011), encoding precision (Bays & Husain, 2008; van den Berg, Awh, & Ma, 2014), false memory and swapping effects (Bays, Catalao, & Husain, 2009), visual long-term memory (VLTM) encoding and retrieval (e.g., Williams, Hong, Kang, Carlisle, & Woodman, 2012; Brady, Konkle, Gill, Oliva, & Alvarez, 2013; Fan & Turk-Browne, 2013), attentional orienting (Golomb, L'Heureux, & Kanwisher, 2014), and ensemble perception (Brady & Alvarez, 2011).

Here we use these tools to examine how VLTM representations and current sensory evidence jointly determine perceptual judgments. As a case study of when these information sources may be most likely to interact, we asked observers to judge the appearance of a familiar object near the threshold of awareness. Specifically, we first associated shapes with unique colors in VLTM, then later presented the shapes again, but now in unrelated colors. Observers estimated the color in which each shape appeared, using continuous report. We manipulated the amount of sensory evidence available during these later presentations by reducing stimulus duration (from 150 ms to 10 ms). We hypothesized that reducing the availability of sensory information would increase reliance on memory, even though long-term memory was entirely task irrelevant (as original and current colors were uncorrelated). By tightly controlling the similarity between the contents

of VLTM and sensory experience (i.e., distance in color space) and the amount of sensory evidence (i.e., duration), this study provides novel quantitative insight into the influence of long-term memory on perception.

Methods

Participants

Across three experiments, 72 observers (49 women, age 18–33 years) who reported normal or corrected-to-normal visual acuity and color vision participated. The sample size was the same for each experiment (24 observers) and was determined based on piloting and previous research. Each observer received course credit or \$12/hr as compensation and provided informed consent to a protocol designed in accordance with the principles expressed in the Declaration of Helsinki and approved by the Institutional Review Board of Princeton University.

Stimuli

Eight “alphabets” of 12 shapes each were used (Fan & Turk-Browne, 2013), for a total of 96 novel items. Each shape repeatedly appeared in the same randomly chosen angular location (radial eccentricity = 8°) and color (CIE L*a*b* space centered at L = 54 / a = 18 / b = -8; radius = 59°). Later, each shape reappeared in the same location but in a new, independently sampled color. Stimuli were presented using a CRT monitor ~70 cm from the observer with MATLAB and PsychToolbox (<http://psychtoolbox.org>).

Procedure

The experiment contained eight blocks, each employing a unique alphabet. Blocks were subdivided into initial exposure and final test phases (Figure 1). For each initial exposure trial, a shape was presented for 500 ms, followed by a 1500-ms blank interval. An achromatic version of the shape (the probe) then reappeared at the same location, along with a color response wheel (radius = 4°). Observers were instructed to recreate the color in which the shape had just appeared by adjusting the probe's color using the color wheel, clicking to report the final selection (Wilken & Ma, 2004). The full alphabet of shapes was repeated three times during initial exposure; the order of shapes within each repetition was randomized. The exposure phase was identical across all three experiments. There were 36 trials/block × 8 blocks = 288 exposure trials in total.




Figure 1. Task display. (a) During initial exposure trials, a colored shape was briefly presented and then an achromatic version of the shape reappeared to probe observers for the color. (b) Final test phase trials were identical, other than the use of forward and backward masks, the duration of the initial shape, and the fact that the shape had previously appeared in a different color.

The final test trials were similar except for three changes: (1) Each shape was presented in a new, randomly chosen color. (2) Shapes were presented between forward and backward color-noise masks of jittered duration (150–225 ms). Noise masks were trial-unique, ring-shaped arrays of randomly colored shape fragments that covered all possible angular locations (inner radius = 6.5° ; outer radius = 9.5°). (3) The shape presentation duration was shortened to 150 ms (Experiment 1), 50 ms (Experiment 2), or 10 ms (Experiment 3). Finally, a random subset of “invisible” trials in each experiment did not contain a shape at all (i.e., 0 ms). Observers were always instructed to report the color of the shape as it appeared on the current trial, identical to the exposure phase, and were informed about “invisible” trials only during debriefing. There were 12 test trials/block \times 8 blocks = 96 test trials in total.

Experiment 1

Twenty-four naïve observers participated (17 women, mean age = 19.3 years). Shapes were assigned in equal proportion to one of two visibility conditions, “150-ms” and “0-ms,” and this assignment was counterbalanced between observers (48 shapes/condition).

Experiment 2

Twenty-four naïve observers participated (15 women, mean age = 21.1 years). One third of all items (32 shapes) were assigned to a 50-ms condition and two thirds (64 shapes) were assigned to a 0-ms condition, to

ensure there was sufficient statistical power in the 0-ms condition to derive robust model parameter estimates. Upon reporting the color on each trial, observers indicated whether they detected the shape using a 4-point scale in response to the question, “How confident are you that you saw the shape?”: *sure present*, *maybe present*, *maybe absent*, or *sure absent*.

Experiment 3

Twenty-four naïve observers participated (17 women, mean age = 22.5 years). In this experiment, two thirds of items (64 shapes) were assigned to a 10-ms condition, and the remaining one third (32 shapes) to a 0-ms condition, to ensure sufficient statistical power in the 10-ms condition to derive robust model parameter estimates. As in Experiment 2, 4-point present/absent responses were collected.

Data Analysis

Because of the small number of final test trials collected per condition per observer, errors were pooled across observers within each condition for subsequent analysis, providing sufficient statistical power to derive robust model estimates. Random-effects reliability was established across observers using bootstrapping (see below).

Modeling performance on initial exposure trials

During initial exposure, there was only one feature associated with each shape. To model performance on




Figure 2. Modeling approach at test. Responses across observers in the final test phase were modeled as a mixture of three component distributions: a uniform distribution to capture guess responses and von Mises distributions to capture responses centered on the “original” color of the shape from exposure and its “current” color, respectively. β = probability of anchoring on a color; σ^{-1} = precision of anchored responses.

these trials, we employed a mixture-modeling procedure that uses maximum-likelihood estimation to partition errors into two component distributions (Equation 1): a uniform distribution reflecting the probability of guessing (inversely related to the proportion of trials on which the response was “anchored” around the true color value); and a von Mises distribution of the anchored responses whose variance provides a measure of the precision (1/SD) of the color representation (Zhang & Luck, 2008).

$$p(\hat{\theta}) = \beta \phi_{\sigma}(\hat{\theta} - \theta) + (1 - \beta) \frac{1}{360} \quad (1)$$

β represents the frequency of responding from the color representation, where ϕ denotes the von Mises distribution centered on the correct feature value θ , with standard deviation σ .

Modeling performance on final test trials

On test trials containing a stimulus (“visible” trials), observers may have used one of two distinct sources of information to generate responses: long-term memory for the color associated with the shape on previous encounters or recent sensory information about the color of the shape on the current trial (Figure 2). To estimate the likelihood of responding based on long-term memory, we adapted the mixture model above to include three components (Equation 2): a von Mises distribution for the “original” color, describing the precision and probability of anchoring upon the long-term memory representation of feature value θ_{orig} from exposure; a von Mises distribution for the “current”

color, describing the precision and probability of anchoring upon the sensory representation of feature value θ_{curr} that just appeared; and a uniform distribution describing the probability of guessing.

$$p(\hat{\theta}) = \beta_{\text{orig}} \phi_{\sigma_{\text{orig}}}(\hat{\theta} - \theta_{\text{orig}}) + \beta_{\text{curr}} \phi_{\sigma_{\text{curr}}}(\hat{\theta} - \theta_{\text{curr}}) + (1 - \beta_{\text{orig}} - \beta_{\text{curr}}) \frac{1}{360} \quad (2)$$

According to this model, perceptual judgment entails sampling directly from *either* long-term memory or sensory representations, resulting in a mixture of response distributions that reflect the independent origins of the underlying signals. This modeling approach has also been used to capture “swap” errors in continuous-report tasks, in which multiple items were concurrently presented but only a single target item was cued for report (Bays et al., 2009). In that case, swaps occurred when a distractor item from the current trial was reported, as opposed to our case, in which swaps reflect a substitution from long-term memory.

Assessing reliability of parameter estimates

As responses were pooled across observers, the above modeling approaches produced single sets of parameter estimates for each condition. To determine the random-effects reliability of these estimates within experiments, we used a bootstrapping procedure (Efron & Tibshirani, 1986). On each of 1,000 iterations, 24 observers’ worth of data were randomly sampled with replacement from the original 24 observers, and the models were refit to this bootstrapped pooled data set.

Experiment	Anchoring probability (β)		Precision (σ^{-1})	
	MLE	95% CI	MLE	95% CI
E1 _(150-ms)	0.962	[0.940, 0.981]	0.074	[0.067, 0.081]
E2 _(50-ms)	0.967	[0.955, 0.978]	0.071	[0.064, 0.078]
E3 _(10-ms)	0.974	[0.964, 0.983]	0.078	[0.072, 0.085]

Table 1. Model parameter estimates for responses during the initial exposure phase.

The logic underlying this resampling approach is that insofar as the effect is reliable across observers, similar results will be obtained regardless of the subset of observers sampled on a given iteration. Across iterations, this procedure yielded an empirical distribution for each model parameter. These distributions were used to construct confidence intervals, as well as to assess model performance. We have used this approach in other published studies that involve single-trial analysis (Kim, Lewis-Peacock, Norman, & Turk-Browne, 2014; Hindy & Turk-Browne, 2016).

Assessing the fit of the three-component model

To evaluate the general validity of the three-component mixture model (Equation 2), we compared its Akaike Information Component (AIC) scores (Akaike, 1974; Burnham & Anderson, 2004) with those of two reduced versions of the model and a null model: (1) a two-parameter version containing only the current-color component (thus ignoring long-term memory), plus guessing; (2) a two-parameter version containing only the original-color component (thus ignoring immediate sensory information), plus guessing; and (3) a zero-parameter null model, containing only the uniform-guessing component.

To assess variability in our group-level estimate of AIC, we used the same bootstrapping procedure as above to construct confidence intervals for the AIC scores of the full, reduced, and null versions of the model in each experiment (Supplemental Figure S1). The p value reported in the results as evidence that observers' performance exceeded that expected for random responding was defined as the number of iterations on which the null model attained a lower AIC (i.e., a better fit) than the full model, and is labeled p_{AIC} . For invisible trials, the full model contained only the original-color component, plus guessing.

Comparing parameters across experiments

To evaluate the consequences of reducing sensory evidence on final test trials, we performed a randomization test across groups that assessed the statistical reliability of trends in model parameters across experiments. Such trends were quantified as the slope of the line

fit to observed parameter values versus stimulus duration (10, 50, 150 ms). A null distribution of slope values was derived by fitting a line 1,000 times to three random, equally sized subdivisions of the set of all observers. Two-tailed p values were derived by evaluating the proportion of values in this distribution that exceeded the observed slope in magnitude. Follow-up analyses of pairwise differences in model parameters between experiments were conducted using a similar procedure. For a given pair of experiments, the null distribution of differences in parameter values was derived by fitting the model 1,000 times to two random, equally sized subdivisions of observers from both groups. Two-tailed p values were derived by looking up the observed difference in this null distribution and identifying the proportion of iterations that exceeded it in magnitude.

Results

Initial exposure trials

We first assessed performance on initial exposure trials to verify successful encoding of color associations into VLTMs. Performance on these trials exceeded chance across all three presentations of each shape (all experiments, $p_{\text{AIC}} \ll 0.001$). These reports were not only statistically reliable but also highly accurate in absolute terms, with anchoring probability exceeding 96% and precision exceeding $0.071\sigma^{-1}$ in all experiments (Table 1).

Visible test trials

The task remained exactly the same on final test trials (Figure 3), except that the shape (when visible) was presented for different durations (across experiments), was visually masked, and was rendered in a new, independently sampled color (θ_{curr}). Performance on these trials greatly exceeded chance (all experiments, $p_{\text{AIC}} < 0.001$). Confirming that the duration manipulation influenced performance, present/absent judgments in Experiment 3 (10 ms) were less sensitive (Table 2; $t = 6.3$, $p < 0.001$) and more conservative ($t = 4.4$, $p < 0.001$) than in Experiment 2 (50 ms).¹

We hypothesized that reducing the amount of sensory evidence for the current color would result in a poorer representation of the current color of the shape. Indeed, as stimulus durations decreased (150 ms to 50 ms to 10 ms), so did anchoring proportion (slope: $p < 0.001$; E1 vs. E2: $p = 0.002$; E2 vs. E3: $p = 0.025$) and precision (slope: $p < 0.001$; E1 vs. E2: $p < 0.001$; E2 vs. E3: $p = 0.003$).

This is consistent with the hypothesis that a decrease in the amount of sensory evidence leads to a




Figure 3. Distribution of errors in test phase. Each panel contains a histogram of errors pooled across all observers and the corresponding fit from the three-component model (dotted line). In (a) and (b), all visible trials are plotted: In (a), errors represent the difference between each response and the current color, whereas in (b), errors represent the difference between each response and the original color. (c) On invisible trials, responses clustered around the original color.

concomitant increase in substitutions from VLTMs (Figure 4). Indeed, the probability that an observer’s report was anchored on the original color increased as stimulus duration decreased (Table 3; slope: $p = 0.003$; E1 vs. E2: $p = 0.006$; E2 vs. E3: $p = 0.075$). Duration did not, however, affect the precision of these reports from long-term memory (slope: $p = 0.215$).

Invisible test trials

No shape was presented at all (invisible or “0-ms” condition) on a subset of final test trials in all experiments, although visual masking disguised these

trials from observers. We included these trials to isolate the contribution of VLTMs to perceptual judgments, as no sensory information was available. Overall, performance on these trials exceeded chance (all experiments, $p_{AIC} < 0.007$). Intrusions occurred both on trials where observers erroneously reported that the shape had been

Experiment	Sensitivity (d')	Criterion
E2 _(50-ms)	2.27	0.341
E3 _(10-ms)	0.753	−0.246

Table 2. Present/absent performance during the final test phase. Response criterion was calculated as in Snodgrass and Corwin (1988).




Figure 4. Modeling results for test phase. (a) The probability of anchoring on the current color and the precision of these responses decreased as stimulus duration was reduced from 150 ms to 50 ms to 10 ms. (b) This reduction increased the probability of anchoring on the original color. At shorter durations (50 ms, 10 ms), the precision of responses anchored on the original color was high, at least relative to the precision for the original color on initial exposure trials (see Table 1). (c) On trials containing no colored shape, anchoring on the original color occurred reliably, and at a consistent rate across experiments, with similar and high precision. Error bars represent bootstrapped 68% confidence intervals (comparable in width to 1 SEM).

present ($p_{AIC} < 0.018$)—thus misattributing their report to sensory input on the current trial—as well as on trials where they correctly reported that the shape had been absent ($p_{AIC} < 0.031$; Table 4). On these “present” trials, anchoring on the original color in VLTM was as likely (Experiment 3; $p = 0.276$) or more likely (Experiment 2; $p = 0.030$) than on “absent” trials.² Neither anchoring probability nor precision differed across experiments (slope: $ps > 0.167$).

Testing an alternative “blend” model

The modeling approach so far assumes that a perceptual judgment entails sampling directly from *either* memory or sensory representations (“swap” mechanism), resulting in a bimodal mixture of response

distributions that reflects the independent origins of the underlying signals. Another possibility, however, is that perceptual judgment entails sampling from an integrated representation (“blend” mechanism) that is formed by combining memory and sensory representations according to Bayes’ rule, resulting a single unimodal response distribution (Knill & Pouget, 2004). Here we evaluate how well such a mechanism might explain our findings.

Formally, this Bayesian blend account asserts that responses are generated by sampling from the posterior density function, which is proportional to the product of three functions: a likelihood function based on the original color seen during exposure trials, a second likelihood function based on current sensory information, and the prior density function representing the probability of the shape taking on any given color.

Experiment	Anchoring probability (β)				Precision (σ^{-1})			
	Current		Original		Current		Original	
	MLE	95% CI	MLE	95% CI	MLE	95% CI	MLE	95% CI
E1 _(150-ms)	0.972	[0.900, 0.990]	0.000	[0.000, 0.032]	0.047	[0.043, 0.051]	0.028	[0.006, 0.080]
E2 _(50-ms)	0.729	[0.616, 0.848]	0.086	[0.021, 0.176]	0.025	[0.020, 0.031]	0.063	[0.022, 0.084]
E3 _(10-ms)	0.251	[0.146, 0.685]	0.216	[0.124, 0.336]	0.017	[0.010, 0.024]	0.065	[0.038, 0.087]

Table 3. Model parameter estimates for the current and original feature values when the stimulus was visible during final test phase.

Experiment	Anchoring probability (β)		Precision (σ^{-1})	
	MLE	95% CI	MLE	95% CI
E1 _(0-ms)	0.144	[0.080, 0.227]	0.060	[0.036, 0.078]
E2 _(0-ms)	0.144	[0.046, 0.263]	0.062	[0.044, 0.086]
“Present” only	0.240	[0.093, 0.382]	0.056	[0.036, 0.115]
“Absent” only	0.119	[0.031, 0.242]	0.065	[0.042, 0.095]
E3 _(0-ms)	0.243	[0.138, 0.384]	0.060	[0.042, 0.084]
“Present” only	0.323	[0.163, 0.499]	0.063	[0.046, 0.102]
“Absent” only	0.216	[0.102, 0.375]	0.057	[0.037, 0.088]

Table 4. Model parameter estimates for responses anchored on the original feature value when the stimulus was invisible during final test phase. Rows labeled “present” contain parameter estimates based only on trials where the observer erroneously reported that the shape was present; “absent” rows were based on trials where the shape was correctly reported as absent.

This latter probability is uniform across colors, so it does not affect the proportionality relationship:

$$p(\hat{\theta}|\theta_{\text{orig}}, \theta_{\text{curr}}) \propto p(\theta_{\text{orig}}|\hat{\theta})p(\theta_{\text{curr}}|\hat{\theta}) \quad (3)$$

We again assume that the current and original color representations follow von Mises distributions, each centered on the correct feature value θ_{curr} and θ_{orig} , respectively, with concentration parameters κ_{curr} and κ_{orig} capturing the amount of uncertainty in the representation of the current color and noise in the long-term memory representation, respectively. The product of these two von Mises distributions is also a von Mises distribution centered on θ_{blend} with concentration parameter κ_{blend} , given by:

$$\theta_{\text{blend}} = \text{atan2} \left(\kappa_{\text{orig}} \sin(\theta_{\text{orig}}) + \kappa_{\text{curr}} \sin(\theta_{\text{curr}}), \kappa_{\text{orig}} \cos(\theta_{\text{orig}}) + \kappa_{\text{curr}} \cos(\theta_{\text{curr}}) \right)$$

$$\kappa_{\text{blend}} = \sqrt{\kappa_{\text{orig}}^2 + \kappa_{\text{curr}}^2 + 2\kappa_{\text{orig}}\kappa_{\text{curr}}\cos(\theta_{\text{curr}} - \theta_{\text{orig}})} \quad (4)$$

We used maximum-likelihood estimation to derive estimates of κ_{curr} and κ_{orig} that provided the best fit to

Experiment	Anchoring probability (β)	Original precision (σ^{-1})	Current precision (σ^{-1})	AIC _{blend}	AIC _{swap}
E1 _(150-ms)	0.965	0.011	0.047	10,451.80	10,446.05
E2 _(50-ms)	0.727	0.011	0.025	8,309.03	8,279.02
E3 _(10-ms)	0.267	0.067	0.014	17,584.33	17,511.62

Table 5. Blend-model parameter estimates for visible trials during the test phase. For side-by-side model comparison, the final column reports the corresponding AIC for the swap model.

errors on visible test trials across all three experiments, then converted these κ values to σ^{-1} values in units of degrees for ease of interpretation. To better account for large errors, a weight parameter β was also estimated, reflecting the contribution of this integrated representation relative to a uniform “guessing” distribution:

$$p(\hat{\theta}) = \beta_{\text{blend}} \phi_{\sigma_{\text{blend}}}(\hat{\theta} - \theta_{\text{blend}}) + (1 - \beta) \frac{1}{360} \quad (5)$$

Interestingly, the resulting parameter estimates for this blend model suggest an overall interpretation that is compatible with those based on the swap model, namely, that the contribution of the current color representation to the posterior diminishes (i.e., its precision decreases) as presentation duration is reduced, whereas the contribution of the original color representation increases (Table 5). These trends are accompanied by increasing guessing proportions, possibly reflecting greater overall response uncertainty at shorter durations. Nevertheless, we found that the blend model consistently underperformed relative to the swap model, as quantified by the difference in AIC (Table 5; Supplemental Figure S1). Thus, although the blend model yields interpretable parameter estimates, the data suggest that the swap mechanism is more likely.

Consistent with these results, visual inspection of normalized error distributions (Figure 5) in the 50-ms and 10-ms conditions suggests that they are bimodal, with modes centered on the current and original colors. This is inconsistent with the blend account, which, critically, predicts a unimodal distribution that peaks between these two color values. Taken together, these findings lend further support to the idea that VLTM manifests in perceptual judgments as discrete substitutions, as expressed by the swap model.

One caveat we would like to add to these model-comparison results is that the AIC differences *per participant* are small (range: 0.24–3.03), as each model was fit to data aggregated across 24 participants. As a rule of thumb, AIC differences smaller than 5 should be interpreted with caution (Burnham & Anderson, 2004). Thus, we consider the blend account to still be very theoretically interesting and potentially relevant. We speculate that blend and swap mechanisms need not be mutually exclusive and that some version of each may dominate in different situations. Specifically, large and




Figure 5. Distribution of responses relative to the current and original colors on visible test trials. Errors were normalized by taking the difference between the response and current color and dividing it by the distance between the original color and current color for each item. This results in a distribution in which the current and original colors have been remapped to 0 and 1, respectively, and errors have been scaled in proportion.

random disparities between old and new information (as in the current study) may promote the formation of more discrete representations, leading to memory substitution errors, whereas an alternative scenario with smaller and/or predictable disparities between old and new information may encourage integration of these two sources of information, leading primarily to memory-driven biases at test (Golomb et al., 2014).

In addition, there exists an interesting intermediate possibility between these strict blend and swap mechanisms. The current formulation of the Bayesian “blend” account assumes that the two likelihood functions reflecting VLTm and sensory information are von Mises distributed, but guessing nevertheless occurs on some proportion of trials. If it were the case that failure to retrieve the color association in VLTm or failure to register current sensory information sometimes occurred with independent probabilities, then the posterior distribution would be formed by combining mixtures of von Mises and uniform distributions by multiplication, according to Bayes’ rule. When the disparity between original and current colors is large enough, the resulting posterior distribution is bimodal and very similar in form to the mixture distribution predicted by the swap model, except that the modes are slightly shifted toward each other. This alternative model is intriguing because it provides a subtly distinct account of the source of the bimodality in response distributions—as arising from direct updating of a single underlying representation that supports two feature values versus arising from sampling a mixture of two representations, each supporting one feature value. However, it is several times more computationally intensive to carry out the necessary simulations to compare these two accounts, because the product of VM+U mixtures does not yield parameters with closed-form definitions. Moreover, the large, random disparities used in the current study may make such small

biases especially difficult to detect. Nevertheless, although the current study design was not optimized to resolve differences between these two accounts, we hope that future investigations will employ designs that are more sensitive to measuring small biases in bimodal distributions.

Discussion

The present study aimed to elucidate how object representations in long-term memory and sensory information about a just-seen object mutually inform perceptual judgments. We examined this question by associating colors with shapes in VLTm and then testing how these memories informed judgments about the color of these shapes when they later appeared again. We found independent contributions of both long-term memory and immediate sensory input on responses, with a greater likelihood of drawing upon VLTm as sensory information decreased. This study provided a strong test of the contribution of VLTm, as these memories were both task-irrelevant (observers were instructed to judge the current color) and orthogonal to the current information (original and current colors were uncorrelated). Interestingly, the precision of these memory substitutions was comparable to the precision with which colors were initially encoded (see also Brady et al., 2013) and was not modulated by the amount of sensory input.

Although we interpret these findings as reflecting automatic substitutions from VLTm during perceptual judgment, an alternative possibility is that observers employed a controlled strategy of defaulting to the original color on certain trials, such as when they did not see the shape. However, this account is insufficient in several ways: (1) The original color was objectively

uncorrelated with the current color across shapes, making it irrelevant to the task of reporting the current color; (2) observers accurately reported the current color when the shape did appear on test trials, showing that they did not assume that the original color was the correct response on these trials; (3) anchoring probabilities on the original color were reliable but notably smaller (8%–22%) than in explicit long-term memory tests in similar paradigms (40%–65%; Brady et al., 2013; Fan & Turk-Browne, 2013), suggesting that anchoring did not simply reflect the accessibility of representations in VLTM; (4) observers reliably responded that a shape had been “present” after reporting the original color on invisible trials, thus misattributing their report to sensory input on the current trial; and (5) on these “present” trials, anchoring on the original color was as likely (Experiment 3; $p = 0.276$) or more likely (Experiment 2; $p = 0.030$) than on “absent” trials.

Together, these points are more consistent with automatic than controlled expression of VLTM during perceptual judgment. Specifically, the representation of the original color is reactivated when there is sensory overlap with specific prior experiences (i.e., a matching shape cue), and this reactivation from memory may be falsely attributed to immediate sensory experience, thereby influencing perceptual judgments even when this influence confers no benefit to (and perhaps even impairs) task performance.

In this study, we collected present/absent judgments to assess the relationship between subjective judgments about the visibility of the current shape and color responses. A promising direction for future research would be to also measure the relationship between subjective confidence in the precision of the current color representation and the rate of substitutions from visual long-term memory, which would provide insight into the metacognitive processes that arbitrate between sensory and long-term representations on a moment-to-moment basis (Suchow, Fougne, & Alvarez, 2012).

Throughout, we employed a mixture-modeling approach to gain basic insights into interactions between long-term memory and sensory information in perceptual judgments. In future studies, alternative modeling techniques, in which encoding precision is modeled as varying across items and trials (Fougne et al., 2012; van den Berg et al., 2012) and/or determined by variability in spiking activity of neural populations (Bays, 2014), may yield a more nuanced understanding of how such interactions influence the content and fidelity of visual representations (see also Supplementary Figure S2).

Taken as a whole, our findings suggest that the visual system automatically recruits long-term memory during perceptual decision making. Ultimately, such retrieval mechanisms may be crucial for deploying past

experience efficiently and for improving perceptual inference.

Keywords: object perception, episodic memory, visual awareness, model fitting

Acknowledgments

We thank Max Luo, Annie Chen, and Jennifer Rawding for help with data collection. We are grateful to Mariam Aly, Aaron Bornstein, Weiji Ma, Jonathan Pillow, Michael Shvartsman, Jordan Suchow, and the Turk-Browne lab for helpful discussions. This work was supported by NSF GRFP DGE-0646086 to J. E. F., NIH F32-EY021999 to J. B. H., and NIH R01-EY021755 to N. B. T.-B.

*Equal contribution.

Commercial relationships: none.

Corresponding author: Judith E. Fan.

Email: jefan@princeton.edu.

Address: Peretsman-Scully Hall, Department of Psychology, Princeton University, Princeton, NJ USA.

Footnotes

¹ These judgments were not collected in Experiment 1.

² These p values reflect the proportion of iterations from the random-effects bootstrapping procedure on which the anchoring probability for “absent” trials exceeded that for “present” trials, multiplied by 2 (two-tailed).

References

- Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, 19, 716–723.
- Ball, K., & Sekuler, R. (1987). Direction-specific improvement in motion discrimination. *Vision Research*, 27, 953–965.
- Bays, P. M. (2014). Noise in neural populations accounts for errors in working memory. *Journal of Neuroscience*, 34, 3632–3645.
- Bays, P. M., Catalao, R. F. G., & Husain, M. (2009). The precision of visual working memory is set by allocation of a shared resource. *Journal of Vision*, 9(10):7, 1–11, doi:10.1167/9.10.7. [PubMed] [Article]

- Bays, P. M., & Husain, M. (2008). Dynamic shifts of limited working memory resources in human vision. *Science*, *321*, 851–854.
- Bays, P. M., Wu, E. Y., & Husain, M. (2011). Storage and binding of object features in visual working memory. *Neuropsychologia*, *49*, 1622–1631.
- Brady, T. F., & Alvarez, G. A. (2011). Hierarchical encoding in visual working memory: Ensemble statistics bias memory for individual items. *Psychological Science*, *22*, 384–392.
- Brady, T. F., Konkle, T., & Alvarez, G. A. (2011). A review of visual memory capacity: Beyond individual items and toward structured representations. *Journal of Vision*, *11*(5):4: 1-34, doi:10.1167/11.5.4. [PubMed] [Article]
- 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.
- Burnham, K. P., & Anderson, D. R. (2004). Multi-model inference: Understanding AIC and BIC in model selection. *Sociological Methods & Research*, *33*, 261–304.
- Cave, C. B. (1997). Very long-lasting priming in picture naming. *Psychological Science*, *8*(4), 322–325.
- Chalk, M., Seitz, A. R., & Series, P. (2010). Rapidly learned stimulus expectations alter perception of motion. *Journal of Vision*, *10*(8):2, 1–18, doi:10.1167/10.8.2. [PubMed] [Article]
- Efron, B., & Tibshirani, R. (1986). Bootstrap methods for standard errors, confidence intervals, and other measures of statistical accuracy. *Statistical Science*, *1*, 54–75.
- Fan, J. E., & Turk-Browne, N. B. (2013). Internal attention to features in visual short-term memory guides object learning. *Cognition*, *129*, 292–308.
- Fougnie, D., Asplund, C. L., & Marois, R. (2010). What are the units of storage in visual working memory? *Journal of Vision*, *10*(12):27, 1–11, doi:10.1167/10.12.27. [PubMed] [Article]
- Fougnie, D., Suchow, J. W., & Alvarez, G. A. (2012). Variability in the quality of visual working memory. *Nature Communications*, *3*, 1229.
- Golomb, J. D., L'Heureux, Z. E., & Kanwisher, N. (2014). Feature-binding errors after eye movements and shifts of attention. *Psychological Science*, *25*, 1067–1078.
- Hansen, T., Olkkonen, M., Walter, S., & Gegenfurtner, K. R. (2006). Memory modulates color appearance. *Nature Neuroscience*, *9*, 1367–1368.
- Hindy, N. C., & Turk-Browne, N. B. (2016). Action-based learning of multistate objects in the medial temporal lobe. *Cerebral Cortex*, *26*, 1853–1865.
- Huang, J., & Sekular, R. (2010). Distortions in recall from visual memory: Two classes of attractors at work. *Journal of Vision*, *10*(2):24, 1–27, doi:10.1167/10.2.24. [PubMed] [Article]
- Jacoby, L. L., Allan, L. G., Collins, J. C., & Larwill, L. K. (1988). Memory influences subjective experience: Noise judgments. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *14*(2), 240.
- Kim, G., Lewis-Peacock, J. A., Norman, K. A., & Turk-Browne, N. B. (2014). Pruning of memories by context-based prediction error. *Proceedings of the National Academy of Sciences, USA*, *111*, 8997–9002.
- Knill, D. C., & Pouget, A. (2004). The Bayesian brain: The role of uncertainty in neural coding and computation. *Trends in Neurosciences*, *27*, 712–719.
- Sagi, D. (2011). Perceptual learning in vision research. *Vision Research*, *51*, 1552–1566.
- Snodgrass, J. G., & Corwin, J. (1988). Pragmatics of measuring recognition memory: Applications to dementia and amnesia. *Journal of Experimental Psychology: General*, *117*, 34–50.
- Suchow, J. W., Brady, T. F., Fougnie, D., & Alvarez, G. A. (2013). Modeling visual working memory with the MemToolbox. *Journal of Vision*, *13*(10):9, 1–8, doi:10.1167/13.10.9. [PubMed] [Article]
- Suchow, J. W., Fougnie, D., & Alvarez, G. A. (2012). Visual working metamemory. *Journal of Vision*, *12*(9): 348, doi:10.1167/12.9.348. [Abstract]
- Tulving, E., & Schacter, D. L. (1990). Priming and human memory systems. *Science*, *247*, 301–306.
- van den Berg, R., Awh, E., & Ma, W. J. (2014). Factorial comparison of working memory models. *Psychological Review*, *121*, 124–149.
- van den Berg, R., Shin, H., Chou, W.-C., George, R., & Ma, W. J. (2012). Variability in encoding precision accounts for visual short-term memory limitations. *Proceedings of the National Academy of Sciences, USA*, *109*, 8780–8785.
- von Helmholtz, H. (1867). *Treatise on physiological optics (Vol. 3)*. Leipzig, Germany: Voss. Retrieved from <http://psych.upenn.edu/backuslab/helmholtz>
- Warrington, E. K., & Weiskrantz, L. (1974). The effect of prior learning on subsequent retention in amnesic patients. *Neuropsychologia*, *12*(4), 419–428.
- Watanabe, T., Náñez, J. E., & Sasaki, Y. (2001). Perceptual learning without perception. *Nature*, *413*, 844–848.

- Wilken, P., & Ma, W. J. (2004). A detection theory account of change detection. *Journal of Vision*, 4(12):11, 1120–1135, doi:10.1167/4.12.11. [PubMed] [Article]
- Williams, M., Hong, S. W., Kang, M.-S., Carlisle, N. B., & Woodman, G. F. (2012). The benefit of forgetting. *Psychonomic Bulletin & Review*, 20, 348–355.
- Witherspoon, D., & Allan, L. G. (1985). The effect of a prior presentation on temporal judgments in a perceptual identification task. *Memory & Cognition*, 13(2), 101–111.
- Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representations in visual working memory. *Nature*, 453, 233–235.