Internal attention to features in visual short-term memory guides object learning

Judith E. Fan and **Nicholas B. Turk-Browne**
Department of Psychology, Princeton University

Abstract

Attending to objects in the world affects how we perceive and remember them. What are the consequences of attending to an object in mind? In particular, how does reporting the features of a recently seen object guide visual learning? In three experiments, observers were presented with abstract shapes in a particular color, orientation, and location. After viewing each object, observers were cued to report one feature from visual short-term memory (VSTM). In a subsequent test, observers were cued to report features of the same objects from visual long-term memory (VLTM). We tested whether reporting a feature from VSTM: (1) enhances VLTM for just that feature (practice-benefit hypothesis), (2) enhances VLTM for all features (object-based hypothesis), or (3) simultaneously enhances VLTM for that feature and suppresses VLTM for unreported features (feature-competition hypothesis). The results provided support for the feature-competition hypothesis, whereby the representation of an object in VLTM was biased towards features reported from VSTM and away from unreported features (Experiment 1). This bias could not be explained by the amount of sensory exposure or response learning (Experiment 2) and was amplified by the reporting of multiple features (Experiment 3). Taken together, these results suggest that selective internal attention induces competitive dynamics among features during visual learning, flexibly tuning object representations to align with prior mnemonic goals.

Keywords

features and objects; memory encoding; memory retrieval; selective attention; visual short-term memory; visual working memory

1. Introduction

Physical actions bring about lasting changes to objects we encounter. These interventions may be more consequential in some cases (e.g., transforming a forest into pulp) than others (e.g., thumbing through a new book), but totally inconsequential actions are rare. Just as physical objects bear traces of such encounters, mental objects may reflect their history of manipulation by the mind. This notion of 'cognitive actions' refers broadly to processes that manipulate mental representations in various ways. Understanding the consequences of such actions is especially important because they are the essence of human cognition (Andrews-Hanna, Reidler, Huang, & Buckner, 2010; Klinger & Cox, 1987; Mason et al., 2007; Singer,

Please address correspondence to either: Judith Fan or Nicholas Turk-Browne, Green Hall, Department of Psychology, Princeton University, Princeton, NJ 08540, jefan@princeton.edu or ntb@princeton.edu, Telephone: (609) 258-4285 Fax: (609) 259-1113.

1966). As a case study of cognitive actions, here we investigate the consequences of reflective, or internal attention (Chun, Golomb, & Turk-Browne, 2011; Chun & Johnson, 2011).

1.1. External and Internal Attention

Attention typically refers to the prioritization of sensory information that is either inherently salient (Theeuwes, 1992; Treisman & Gelade, 1980) or relevant to current goals (Most, Scholl, Clifford, & Simons, 2005). Deploying attention to certain information and not other information does not merely determine what gets processed downstream (Al-Aidroos, Said, & Turk-Browne, 2012), but can also alter the perceptual experience of this information (Carrasco, Ling, & Read, 2004). There are both benefits and costs of attention: processing of selected information can be facilitated (Posner, Snyder, & Davidson, 1980) and processing of unselected information can be inhibited (Houghton & Tipper, 1994).

These forms of attentional modulation can have longer-term consequences for memory. On the one hand, allocating attention to an item during encoding enhances its later recognition (Chun & Turk-Browne, 2007; Uncapher, Hutchinson, & Wagner, 2011). On the other hand, removing attention from an item during encoding not only worsens recognition, but can actually produce a memory cost for this item relative to novel items (Fox, 1995; Tipper, 1985). This cost may reflect an attentional weighting mechanism that actively inhibits distractors, shaping the representation of an item in memory based on current task goals (Lavie & Fox, 2000; Neill & Valdes, 1992).

The impact of attention on memory has been investigated mostly in the case of external attention, which entails the selection of representations that are directly supported by sensory information, and that do not require maintenance in working memory to remain active. Here we investigate the mnemonic consequences of *internal* attention, which we define as the selection of representations that are being maintained temporarily in working memory, and that are no longer supported by externally available sensory information. External attention might nevertheless provide a useful starting point for thinking about how internal attention affects memory. This analogy is supported by the fact that external and internal attention rely on similar mechanisms: (1) they are both highly selective and capacity limited (Chun et al., 2011), (2) deploying one form of attention interferes with the other (Awh & Jonides, 2001), and (3) both forms of attention engage overlapping networks of brain regions (Nobre et al., 2004).

1.2 Consequences of Memory Retrieval

Given these commonalities to external attention, memory may be enhanced when internal attention is allocated to an item and suppressed when it is removed from an item. Support for this prediction comes from research on retrieval from long-term memory. While overly simple conceptions of remembering liken retrieval to finding and reading a file from a hard drive, the act of recovering the contents of a representation from long-term memory actually alters these contents and affects subsequent retrieval of the same information.

These alterations can be beneficial. For example, testing of recently learned information—such as a history lesson (Nungester & Duchastel, 1982) or word pairs (Carpenter, Pashler, & Vul, 2006; Carrier & Pashler, 2002)—promotes long-term retention of this information. Moreover, retention is better after repeated testing (i.e., after multiple retrieval attempts) than after an equal number of opportunities to study the same material (Karpicke & Roediger, 2008; Roediger & Karpicke, 2006). Relatedly, information tends to be better remembered when it has been generated from internal knowledge than when acquired via external sources (Crutcher & Healy, 1989; Slamecka & Graf, 1978).

Memory retrieval can also have deleterious effects. For example, recall of one item from memory makes other related items less accessible for later recall (i.e., retrieval induced forgetting; (Anderson, Bjork, & Bjork, 1994; 2000). Specifically, in such studies, participants study a list of category-exemplar word pairs (e.g., Fruit-Grape, Fruit-Peach, Mammal-Cow), practice recalling a subset of the exemplars with a category cue (e.g., Fruit-G___?), and finally, perform a recollection test of all exemplars. Memory is better for the pairs practiced in the second phase (e.g., Fruit-Grape) relative to unpracticed pairs (e.g., Mammal-Cow), consistent with the testing effects described above. Among unpracticed pairs, however, memory is worse for those exemplars that share a category cue with a practiced exemplar (e.g., Fruit-Peach). These results are interpreted as evidence that inhibitory mechanisms suppress memory of competitors during initial recall.

1.3 Internal Attention and Visual Learning

Previous work has focused on the retrieval of stable information from long-term memory, such as semantic categories and their members. However, similar competitive processes may operate in a short-term store of recently experienced visual information, when a subset of this information is selected via internal attention. This may help shape how items are represented during initial encoding, and ultimately guide longer-term learning about these items. This view is compatible with modal models of memory, which emphasize the importance of a short-term store as the nexus between ongoing perception and long-term knowledge (Atkinson & Shiffrin, 1969). Indeed, the hallmark of short-term memory is that representations can be manipulated in the service of ongoing behavior—including learning (Baddeley and Hitch, 1974).

Visual memory provides a fruitful domain in which to investigate the impact of internal attention on long-term learning. Specifically, what constitutes a unit of visual memory remains contentious (Fougnie, Asplund, & Marois, 2010), including theories based on: objects (Luck & Vogel, 1997), feature dimensions (Olson & Jiang, 2002; Wheeler & Treisman, 2002), and information load (Alvarez & Cavanagh, 2004; Bays & Husain, 2008; Brady, Konkle, & Alvarez, 2009). All of these theories share an assumption that the contents of memory are solely determined by properties of stimuli in the current display. However, we propose that the lack of consensus partly results from neglecting the role of prior experience in shaping how a given display or object is construed. In particular, beyond examining how prior perceptual experience affects visual memory (Brady et al., 2009; Curby, Glazek, & Gauthier, 2009), we consider how prior *retrieval* experience tunes object representations.

Initial insights can be gleaned from a version of the standard retrieval-induced forgetting task that used visual stimuli (Ciranni & Shimamura, 1999). Instead of relying on semantic associations between items, this study manipulated the perceptual similarity of items. Participants learned the location of twelve objects that were grouped by shape (e.g., circles, triangles, crosses) or color (e.g., green, orange, purple), practiced retrieving an associated feature of a subset of items, and were then worse at recalling features of unpracticed items that were perceptually grouped with practiced items. These findings suggest the presence of competition between objects in visual memory.

1.4 The Current Study

Here we investigate the impact of internal attention on object learning by examining the role of selective retrieval of features from visual short-term memory (VSTM). After viewing a complex object defined along multiple feature dimensions (e.g., color, orientation, location), we asked: How does selection of one feature from its representation in short-term memory affect how this object is represented in long-term memory? We aim to address two key

aspects of visual object learning in the current study: First, motivated by the modal model of memory, we test how accessing an object in a short-term store influences how that object comes to be represented in a long-term store. Second, building on extant theories of visual memory, we explore how retrieval history influences the weighting of features within individual objects. These goals are distinct from prior work on retrieval-induced forgetting, for example, which examined the mnemonic consequences of retrieving consolidated representations in long-term memory rather than nascent representations in short-term memory, and which examined competition between items that are semantically or visually related rather than between features within items.

To accomplish these goals, we designed a visual memory task that manipulates feature retrieval history on an item-level basis (Figure 1). Each novel object was characterized by a unique: angular location relative to central fixation, orientation defined as degree of rotation about center of the shape, and color chosen from a perceptually uniform space. On each trial of a practice phase, observers viewed an object in the periphery of the display. After offset, observers were prompted by a post-cue to report one of its features from VSTM by adjusting a memory probe; a post-cue was used to manipulate internal, rather than external, attention. Observers adjusted the probe by moving the mouse cursor along the wheel until it matched the original object on the cued dimension. All objects were viewed an equal number of times during the practice phase.

To measure the strength of visual long-term memory (VLTM) for the features of each object, observers completed a final test phase with only a feature cue and memory probe. An advantage of using location, orientation, and color as features is that these dimensions are continuous and can be mapped onto a circular space. This approach affords precise measurement of the fidelity of a feature representation, by computing the angular deviation of responses from the true value on the circle (Zhang & Luck, 2008).

We conducted three experiments with different groups of observers. In Experiment 1, we explored the relationship between repeated retrieval of a novel object's features from VSTM and subsequent VLTM for this object. Experiment 2 eliminated alternative explanations of our results based on the amount of sensory experience and the possibility for motor learning. Experiment 3 explored the impact of retrieving multiple features from VSTM on VLTM. Together, these studies reveal important consequences of selective internal attention to features on the organization of memory for objects.

2. Experiment 1: Retrieval History and the Tuning of Object Representations

This experiment serves as an initial exploration of the benefits and costs of retrieval from VSTM on VLTM. Objects were assigned to one of three conditions (Figure 2A): In the Constant and Switch conditions, the same feature was probed throughout the practice phase; in the test phase, the same feature (Constant) or an unpracticed feature (Switch) was probed. In the Passive condition, objects were studied the same number of times in the practice phase, but no features were reported; in the test phase, an unpracticed features was probed, providing a baseline measure of feature memory in the absence of retrieval history.

If selection of one feature of an object from VSTM just enhances the representation of that feature in VLTM (practice-benefit hypothesis), then test accuracy should be higher for Constant than Passive, and Switch and Passive should not differ. Alternatively, if selecting a feature enhances the representation of all features (object-based hypothesis), then both Constant and Switch should have higher accuracy than Passive. Finally, if this selective operation induces competition among features vying for retrieval (feature-competition

hypothesis), then test accuracy should be higher for Constant than Passive and lower for Switch than Passive. If VLTM is unaffected by accessing VSTM, all conditions should produce quantitatively similar performance; note that this null hypothesis does not only rely on null effects, as test performance in all conditions should be better than chance.

2.1 Methods

2.1.1 Participants—Thirty-six naïve observers (23 women, mean age 19.8y) participated in this experiment. In all experiments, observers received course credit or $12 for participating, provided informed consent, and reported normal or corrected-to-normal visual acuity and color vision. The study protocol was approved by the Princeton IRB.

2.1.2 Stimuli—Eight 'alphabets' containing eighteen shapes each were used. Shapes were selected from an online repository of non-Roman ideograms (http://symbols.com, HME Publishing) and reproduced freehand in a vector graphics drawing program. Shapes in an alphabet tended to share common contours, but were otherwise perceptually distinct, as well as rotationally and reflectively asymmetrical. We chose these shapes as stimuli because they were highly novel, preventing prior experience from contaminating our investigations of object learning. Each shape served as the base for one object, with which values from three clearly defined feature dimensions (location, orientation, and color) were associated to form more complex objects. This general approach has been used previously in the study of how objects are represented in VSTM (e.g., Luck & Vogel, 1997). Each shape was displayed at a fixed eccentricity (8°), and was assigned a single angular location (relative to central fixation), orientation (degree of rotation about center of shape relative to arbitrary canonical orientation), and color (from a perceptually uniform space: equiluminant CIE L*a*b* color space centered at L = 70, a = 20, b = 38; radius 60). Feature values in each dimension were independently and randomly sampled from 180 circular coordinates. Within each alphabet, no two objects could exhibit the same combination of feature values. Stimuli were presented using MATLAB and PsychToolbox (Brainard 1997; Pelli 1997).

2.1.3 Procedure—The experiment consisted of eight blocks, and each block was randomly assigned a unique stimulus alphabet. Within each block, observers completed two phases: a VSTM practice phase followed by a VLTM test phase. Before starting the experiment, observers were briefed on the two-phase structure of each block and instructed to aim for high accuracy in both practice and test phases of all blocks. A repeated practice-test block design was used instead of a single practice-test session to minimize fatigue and proactive interference during test. That is, while a longer interval between study and test should not have affected VLTM much, we were concerned that these other sources of noise might contaminate our primary dependent measure.

Within a block, each of the eighteen objects in the alphabet was presented three times during the practice phase and once during the test phase. The order in which objects appeared was randomized, with the constraint that in the practice phase, every object had to appear at least once before any other object could repeat. Six objects were assigned to each of the Constant, Switch, and Passive conditions within-subject, and these assignments were counterbalanced across subjects. In the Constant and Switch conditions, the same feature was probed three times in the practice phase, and either this practiced feature (Constant) or an unpracticed feature (Switch) was probed in the test phase. In the Passive condition, objects were presented three times but no features were reported in the practice phase, and an unpracticed feature was probed in the test phase. Across the objects in every condition, the assignment of dimensions to be probed in the practice and test phases was carefully counterbalanced. In each block there were 54 trials in the practice phase (6 objects × 3 conditions × 3 presentations) and 18 trials in the test phase (6 objects × 3 conditions × 1 presentation).

Across blocks, there were a total of 432 VSTM practice trials and 144 VLTM test trials. The full procedure lasted less than an hour with instructions and debriefing.

On each trial of the practice phase, observers were presented with a single object for 1000 ms under free-viewing conditions. Following its offset and a retention interval lasting 1500 ms, observers were prompted by a text post-cue to report one of the object's features from VSTM. Observers freely adjusted this feature of the memory probe—an initially black, canonically oriented, and centered version of the shape—until it matched the original object on the cued dimension. A continuous response wheel surrounded the memory probe, which observers clicked on to select a value: for the color cue, this entailed choosing from a color wheel; for the orientation cue, this entailed clicking the response wheel at the angle that reflected the desired rotation; for the location cue, this entailed clicking the spatial position of the original stimulus. To make Passive trials as similar as possible to the Constant and Switch conditions (which required responses), observers were presented with the memory probe and a blank response wheel, and were instructed to click on the probe itself to proceed. On trials requiring reporting of a feature (Constant/Switch), the mouse cursor was initially positioned at the center of the display, and the observers moved it to the wheel and clicked. On trials not requiring a report (Passive), the mouse cursor was initially positioned at a random angle on the response wheel, and observers moved it to the center and clicked. This central movement in the Passive condition was chosen to ensure that observers attended to the memory probe, which they could have otherwise ignored on these trials. Nevertheless, the magnitude of the movement was equated across conditions. On test phase trials, only the text cue and memory probe were presented, requiring observers to report cued features purely from VLTM.

Presenting one object per trial helped limit possible feature interactions to those between features within an object. A downside of this approach is that observers might have had enough spare resources to represent each feature separately rather than as part of a bound object. However, we felt that the alternative—presenting multiple objects per trial—would have been less ideal for three reasons: (1) simultaneous presentation of multiple objects would increase the likelihood of interactions between features from different objects, obscuring interactions between features from the same object; (2) the binding of features into a coherent object requires focused attention (e.g., Treisman, 2006), and a set size of one ensured that the object would be fully attended; and (3) as set size increases, observers may default to summary statistics over feature dimensions, which can bias individual object representations (Brady & Alvarez, 2011). Nevertheless, future research should examine the consequences of selective feature retrieval from VSTM for larger set sizes.

2.1.4 Measurement of memory performance—Response error was calculated as the angular deviation from the position on the wheel corresponding to the correct value on the cued dimension. Because responses were continuous along each feature dimension, this measure permitted quantitative estimates of the error distributions for each condition. Narrower distributions reflected more accurate performance, with more responses clustered about the true value. To quantify accuracy in each condition, the root average squared deviation from the true value in degrees was analyzed (root mean squared error, RMSE).

RMSE is a standard estimation approach that aggregates errors across individual observations into a single measure of performance and is intuitively expressed in the same units being estimated (in this case, degrees). When observers' responses are unbiased, RMSE is formally equivalent to the standard deviation of errors. There are other prominent approaches for characterizing performance in continuous report tasks, including fitting a mixture of theoretical distributions using maximum likelihood estimation (e.g., Zhang & Luck, 2008; see also Wilken & Ma, 2004). Although we employ this mixture modeling

technique later to explore different components of performance related to "accessibility" and "precision," we focus primarily on RMSE because it provides a more direct measure of performance, requiring fewer assumptions about the underlying response distributions and coping better with the relatively small number of trials we could obtain in the VLTM test.

We modeled chance performance as a uniform distribution over the interval [-180°, 180°) corresponding to the range of possible errors on the report task, where 0° represents the correct value. The expected magnitude of errors made by an observer guessing randomly and with no information about the correct feature value would be 90°. However, the RMSE under pure guessing conditions is given by the expression $\frac{360°}{\sqrt{12}} \approx 103.9°$ (Rinaman, Heil, Strauss, Mascagni, & Sousa, 2012), derived from the definition of standard deviation of a uniform distribution (see dotted line in Figure 3B). The height of its probability density function is $\frac{1}{360°} = 0.0028$ (see dotted line in Figure 3A).

In analyzing RMSE, we collapsed over the feature dimension being probed within each condition. Because the design was carefully counterbalanced such that each feature was equally likely to be practiced and tested, there were many feature permutations and thus we had insufficient statistical power for examining features separately (some permutations occurred as rarely as once per block). Whether there are feature-specific differences in the consequences of retrieval practice remains an important question for future research. This could be examined by reducing the number of dimensions or probing a single dimension for all objects.

Because we were interested in measuring the consequences of retrieval from VSTM on VLTM, only objects whose features were successfully reported during the practice phase were retained for analysis. Objects were excluded from analysis if, on any of the practice phase trials, the RMSE fell 1.5 times the interquartile range (IQR) outside the 1st and 3rd quartiles in each condition (Frigge, Hoaglin, & Iglewicz, 1989; Tukey, 1977). This outlier exclusion procedure was independent of VLTM test performance, our primary dependent measure. On average, 1.02 objects per condition and block were excluded from further analyses as outliers; this rate did not differ between the Constant and Switch conditions ($t(35) = 1.25$, $p = 0.219$, $d = 0.209$).

2.2 Results and Discussion

In the practice phase, performance was better than chance at all three repetitions in both the Constant and Switch conditions (all $t(35)$s > 225.8, ps 0.001, ds > 37.6). We analyzed the practice phase data using a 2 (condition) × 3 (repetition) repeated-measures ANOVA. Differences in accuracy between Constant and Switch were not expected because these conditions were identical until the test phase (Constant RMSE [first, second, third] = 11.35°, 11.05°, 10.63°; Switch RMSE [first, second, third] = 11.11°, 10.83°, 10.72°). Indeed, there was no main effect of condition ($F(1,35) = 0.152$, $p = 0.699$, $\eta_p^2 = 0.004$). Despite high overall accuracy, there was a main effect of repetition with performance improving over the three repetitions ($F(2,70) = 5.51$, $p = 0.006$, $\eta_p^2 = 0.136$). There was no interaction between condition and repetition ($F(2,70) = 0.382$, $p = 0.684$, $\eta_p^2 = 0.011$).

We hypothesized that having reported one feature of an object during the practice phase would, in the test phase, result in better VLTM for that feature (Constant) and worse VLTM for other features of the same object (Switch), relative to baseline (Passive). As can be seen in Figure 3A, aggregate test phase errors were consistent with this prediction (Constant RMSE = 55.60°; Switch RMSE = 75.22°; Passive RMSE = 70.37°). To assess the reliability of these differences between conditions, we compared RMSE across observers (Figure 3B).

Error was reliably lower for Constant than Passive ($t(35) = 6.31$, $p < 0.001$, $d = 1.05$) and higher for Switch than Passive ($t(35) = 2.24$, $p = 0.031$, $d = 0.374$). This latter difference is striking because both Switch and Passive test trials probed unpracticed features, differing only with respect to retrieval history of other features. Performance was better than chance in all three conditions (all $t(35)s > 10.4$, p 0.001, $ds > 1.73$). These results suggest that selecting one feature of a novel object from VSTM biases the resultant representation towards the attended feature and away from unattended features.

A potential explanation for why VLTM performance was worse in the Switch condition is that there was response-level interference between the practiced feature and unpracticed features, especially for features with a spatial component (e.g., location and orientation). For instance, making a motor movement to report the color of an object might interfere more with memory for the object's location than reporting the location of an object would interfere with memory for the object's color. A key prediction of this account is that the degree of VLTM impairment for an unpracticed feature (i.e., the cost for Switch relative to Passive) will differ between objects where a non-spatial feature (i.e., color) was practiced but a spatial feature (i.e., orientation or location) was tested vs. where a spatial feature was practiced but a non-spatial feature was tested. To examine this possibility, we performed a 2 (feature type: spatial vs. non-spatial) $\times$ 2 (condition: Switch vs. Passive) repeated-measures ANOVA on VLTM performance. The main effect of condition remained, with Switch impaired relative to Passive ($F(1,35) = 4.22$, $p = 0.0475$, $\eta_p^2 = 0.108$; p-value differs from above because only a subset of the feature transitions were included). However, inconsistent with the possibility that unpracticed spatial features suffered more interference, there was no interaction between feature type and condition ($F(1,35) = 0.0173$, $p = 0.896$, $\eta_p^2 < 0.001$).

Directing internal attention to information in short-term memory has been shown to enhance subsequent memory (Johnson, Reeder, Raye, & Mitchell, 2002). Our results build on these findings by showing that diverting internal attention away from information can also impair its representation in long-term memory. That is, we interpret our results as evidence that deploying internal attention concurrently enhances selected features while suppressing other unselected features. This tuning may result from an inhibitory mechanism needed to resolve competition between representations of the cued and uncued features of the probe that are simultaneously active in VSTM. This kind of explanation has been applied to memory for arbitrary lexical and semantic associations (Carpenter et al., 2006; Carrier & Pashler, 2002; Levy & Anderson, 2002), but here we demonstrate a novel role for feature-level inhibition during object learning.

3. Experiment 2: Ruling Out Alternative Explanations

This experiment seeks to rule out two alternative explanations for the difference between Switch and Passive in Experiment 1. First, although Switch and Passive conditions were equated in testing an unpracticed feature, they differed in that observers received extra sensory exposure to the cued feature of Switch objects while it was reported in the practice phase. That is, additional encoding (rather than retrieval) of the practiced feature may have strengthened the weight of that feature relative to the unpracticed test feature. Second, observers may have associated Switch objects with the location on the response wheel that was needed to report the cued feature in the practice phase. Because the same wheel was used for all feature dimensions, such response learning could have caused interference when the unpracticed test feature required a different response. A related issue is that the presence of these spatial landmarks and repeated motor trajectories may have inherently interfered with the VSTM representation of unpracticed features with a spatial component, such as location and orientation.

To remove these potential confounds, we modified the task from Experiment 1 by changing how the memory probe was presented and how observers made responses. To equate sensory exposure across dimensions, the memory probe—rather than being black, canonically oriented, and centered—was identical to the original object *except* for the cued dimension, which was set to a random starting value. For example, if post-cued to report color, the memory probe appeared in the same location and orientation as the object, but in a random color. To eliminate the possibility of response learning, the mouse cursor and wheel were removed from the display. (Responses were tracked on a hidden wheel, but observers could only base their response on changes in the memory probe.) Moreover, the mapping between cursor position and feature space was randomly rotated on every trial to ensure that novel mouse trajectories were always required to arrive at any given feature value.

If the results of Experiment 1 reflect response learning, then differences between conditions should be eliminated by the lack of a consistent response in the practice phase. If the results reflect biased sensory exposure, then accuracy should be markedly different than Experiment 1, with Switch similar to Passive, and both better than Constant. Consider the example object above: While reporting color in the practice phase, there was *greater* exposure to the unpracticed location and orientation features because these values were always veridical in the probe; in contrast, the practiced color feature was initially random and only close to the "true" value near the time of response. Thus, if this object was in Constant condition, the amount of exposure to the test feature (color) would be less than if it was in the Switch (veridical location or orientation) or Passive (all features veridical). Alternatively, according to our hypothesis that selecting a feature in VSTM suppresses other unpracticed features in VLTM, we should replicate the lower accuracy for Switch relative to Passive despite equivalent sensory exposure to unpracticed features in both conditions. We did not have strong predictions for Passive vs. Constant, since a benefit of repeated testing for Constant could be diluted by a benefit of greater sensory exposure for Passive.

3.1 Methods

3.1.1 Participants—Thirty naïve observers (17 women, mean age 20.3y) participated in this experiment.

3.1.2 Procedure—The stimuli and procedure were identical to Experiment 1, with four changes (Figure 1B): (1) observers adjusted the mouse position to report the cued feature of the memory probe, but other visual feedback was eliminated (i.e., no response wheel or mouse cursor); (2) the mapping between the (hidden) cursor position and feature space was randomly rotated from trial-to-trial to ensure a unique mouse trajectory each time an object was probed; (3) the memory probe was identical to the original object in appearance, except for the cued dimension, which was set to a random starting value on each trial; (4) in the Passive condition, where no dimension was cued in the practice phase, the original object was presented intact for an interval that matched the running mean response time of all Constant and Switch trials in that block—this delay was used to ensure equal average sensory exposure to all objects.[1] Because this experiment lasted longer, observers completed five practice-test blocks, with alphabets sampled randomly from the original corpus of eight alphabets.

---

[1]These changes were made to address possible alternative explanations related to sensory exposure and motor learning, but they have other nice properties as well. For instance, the task in Experiment 1 may have encouraged observers to treat individual features as separate objects because the memory probe involved reporting one feature in the absence of the other associated features. Here, in contrast, reporting the probed feature required the 'reconstruction' of the full original object and may thus have encouraged more wholistic object-based processing.

3.2 Results and Discussion

In the practice phase, performance was better than chance at all three repetitions in both the Constant and Switch conditions (all $t(29)s > 151.0$, ps 0.001, $ds > 27.6$). On average, 1.19 objects per condition and block were excluded from further analyses as outliers; this rate did not differ between Constant and Switch ($t(29) = 0.359$, $p = 0.722$, $d = 0.0656$). We analyzed the practice phase data using a 2 (condition) × 3 (repetition) repeated-measures ANOVA. Accuracy did not differ between Constant and Switch ($F(1,29) = 0.066$, $p = 0.799$, $\eta_p^2 = 0.002$). Unlike Experiment 1, the improvement across repetitions did not reach significance ($F(2,58) = 2.10$, $p = 0.132$, $\eta_p^2 = 0.067$), although the range of performance was comparable (Constant RMSE [first, second, third] = 10.68°, 10.04°, 10.08°; Switch RMSE [first, second, third] = 10.71°, 10.28°, 10.19°). This suggests that observers may have been performing near ceiling in the VSTM task in this experiment. Regardless, retrieval practice can affect later test memory even when performance is at ceiling (Roediger & Karpicke, 2006). There was no interaction between condition and repetition ($F(2,58) = 0.074$, $p = 0.929$, $\eta_p^2 = 0.003$).

Consistent with our hypothesis, the distribution of errors pooled across observers mirrored that of Experiment 1 (see Figure 4A; Constant RMSE = 56.77°; Passive RMSE = 59.69°; Switch RMSE = 68.86°). Assessing the reliability of these differences across observers (Figure 4B): the decrement for Switch vs. Passive again reached significance ($t(29) = 2.77$, $p = 0.0096$, $d = 0.506$), but the advantage for Constant vs. Passive did not ($t(29) = 1.00$, $p = 0.326$, $d = 0.182$); Constant was nevertheless more accurate than Switch ($t(29) = 5.19$, $p < 0.001$, $d = 0.948$). Finally, performance was better than chance in all conditions (all $t(29)s > 10.8$, ps 0.001, $ds > 1.91$). We placed stress on the hypothesized role of selective feature retrieval in tuning memory for objects by controlling the amount of sensory exposure and eliminating the possibility of response learning. The results suggest that suppression of VLTM for unpracticed features is a robust consequence of selection within VSTM, even when those features are externally available during this internal selection.

Although the overall pattern of results matched Experiment 1, the benefit of retrieval practice for Constant vs. Passive was mitigated. We interpret this smaller effect as the canceling out of a practice benefit for Constant (that we observed in Experiment 1) by additional encoding for Passive (which was not part of Experiment 1). That is, in the Passive condition, all three features were fully present on the memory probe, whereas in the Constant condition, observers needed to regenerate the practiced and later tested feature, moving through intermediate values that may have diluted the representation. Consistent with this interpretation, accuracy in the Constant condition of both experiments was nearly identical (55.60° in Experiment 1 vs. 56.77° in Experiment 2; $t(64) = 0.30$, $p = 0.767$, $d = 0.0737$). What differed between experiments was the Passive condition (70.37° in Experiment 1 vs. 59.69° in Experiment 2; $t(64) = 2.18$, $p = 0.033$, $d = 0.538$), as would be expected if the smaller effect was caused by greater sensory exposure (which only affected unpracticed features). Nevertheless, if the enhancement effect were stronger, it would have overcome this effect of sensory exposure and revealed itself again. For now, the enhancement effect in Experiment 1, while robust, awaits further study and validation (see also Section 5 for additional tentative evidence from estimation of guessing rates).

Another interesting trend emerged when comparing Experiments 1 and 2. Namely, in addition to the reliable difference between Switch and Passive, Switch did not benefit as much from greater sensory exposure as Passive (75.22° in Experiment 1 vs. 68.86° in Experiment 2; $t(64) = 1.51$, $p = 0.136$, $d = 0.373$). In fact, the size of the Switch vs. Passive difference roughly doubled in magnitude. Although this is a post-hoc finding that requires further investigation, one intriguing possibility is that the correspondence between the externally available features of the probe and the representation of the object in VSTM may

have externalized and strengthened the competition—otherwise internally mediated—between active feature representations, leading to greater weakening. Indeed, stimuli that match the contents of working memory automatically capture attention (Soto, Hodsoll, Rotshtein, & Humphreys, 2008). Relatedly, the external availability of unpracticed features may have allowed them to seep into VSTM, resulting in even greater inhibition when the cued feature was attended and reported.

4. Experiment 3: Variable vs. Focused Retrieval History

In Experiments 1 and 2, we examined the impact of selective retrieval of one feature from VSTM on object learning. However, it is not uncommon for more than one feature of an object to be relevant over repeated encounters. This experiment tests how selection of multiple features from VSTM affects VLTM for novel objects.

Objects were assigned to one of three conditions (Figure 2B): In the Constant and Switch conditions, the same feature was probed twice in the practice phase; in the test phase, the same feature or an unpracticed feature was probed, respectively. In the Variable condition, two different features were probed in the practice phase; in the test phase, the third (unpracticed) feature was probed. This design affords a direct comparison between variable practice (Variable) and focused practice (Switch) with respect to long-term memory for unpracticed features.

There were three possible outcomes: First, variable practice might moderate the suppression of unpracticed features by releasing inhibition when previously unpracticed features must be practiced, resulting in higher accuracy in Variable than Switch. Second, the distinction between variable and focused practice may be inconsequential insofar as all that matters for unpracticed features is that they are not selected, in which case Variable and Switch accuracy would not differ (but should be better than chance). Third, variable practice might induce greater suppression of always unpracticed features due to the weak status of previously unpracticed features that must be practiced, which incites strong competition and results in higher accuracy for Switch than Variable.

4.1 Methods

4.1.1 Participants—Thirty naïve observers (18 women, mean age 20.0y) participated in this experiment.

4.1.2 Procedure—All procedures and stimuli were identical to Experiment 2, except that each object was presented twice during the practice phase and the Passive condition was replaced by a Variable condition in which two different features were practiced and the third (unpracticed) feature was tested.

4.2 Results and Discussion

In the practice phase, performance was better than chance for both repetitions in the Constant, Switch, and Variable conditions (all $t(29)s > 138.2$, ps 0.001, $ds > 25.2$). On average, 0.90 objects per condition and block were excluded from further analyses as outliers; this rate did not differ between Constant and Switch ($t(29) = 1.66$, $p = 0.107$, $d = 0.304$) or between Switch and Variable ($t(29) = 1.30$, $p = 0.204$, $d = 0.237$). However, 0.23 more objects were excluded per block from Variable than Constant ($t(29) = 3.12$, $p = 0.0041$, $d = 0.569$). This difference may be attributable to the fact that the Variable condition involved practicing two features. However, since Constant and Switch conditions were identically structured during this phase, and the Switch exclusion rate did not differ from that of Variable, the Constant/Variable difference may not be meaningful. Regardless,

additional analyses verified that the exclusion procedure did not affect the pattern of VLTM results. We analyzed the practice phase data using a 3 (condition) × 2 (repetition) repeated-measures ANOVA. Performance did not differ across conditions ($F(2,58) = 1.01$, $p = 0.372$, $\eta_p^2 = 0.034$), despite the fact that in the second presentation the same feature was practiced for Constant and Switch and a new feature was practiced for Variable. Accuracy did not differ across the two repetitions ($F(1,29) = 1.08$, $p = 0.308$, $\eta_p^2 = 0.036$; Constant RMSE [first, second] = 11.13°, 10.54°; Switch RMSE [first, second] = 10.37°, 10.37°; Variable RMSE [first, second] = 10.91°, 10.96°). There was no interaction between condition and repetition ($F(2,58) = 0.788$, $p = 0.460$, $\eta_p^2 = 0.026$).

Consistent with the other experiments (Figure 5A), VLTM was better for practiced features (Constant RMSE = 60.10°) than unpracticed features in cases where one other feature had been practiced (Switch RMSE = 68.35°). Practicing two features led to even worse VLTM for unpracticed features (Variable RMSE = 75.07°). Assessing the reliability of these differences across observers (Figure 5B): the decrement for Switch vs. Constant reached significance ($t(29) = 2.77$, $p = 0.0096$, $d = 0.506$), as did the decrement for Variable vs. Switch ($t(29) = 3.08$, $p = 0.0045$, $d = 0.563$). Finally, performance was better than chance in all conditions (all $t(29)s > 9.66$, ps ≤ 0.001, ds > 1.76).

Despite equating study and retrieval opportunities, selecting multiple features of an object from VSTM suppressed VLTM for other features more than selecting a single feature repeatedly. Indeed, under both variable and focused practice, an unpracticed feature was finally tested—the objects differed only with respect to the nature of prior retrieval practice. These findings contradict the simplest conception of retrieval tuning, whereby only the number of retrieval attempts predicts the strength of subsequent memory for practiced and unpracticed features. Moreover, lower accuracy under variable practice stands in contrast to other examples of beneficial variability, such as 'double-training' in perceptual learning (Xiao et al., 2008) and the 'spacing effect' in episodic memory (Melton, 1970). Instead, these results are reminiscent of the finding that 'extinction' is stronger under variability: When a conditioned stimulus is not reinforced in multiple contexts vs. a single context, then subsequent renewal of its association with an unconditioned stimulus is attenuated (Chelonis, Calton, Hart, & Schachtman, 1999). Although a speculative connection at this stage, to the extent that the color, orientation, and location reports served as distinct task contexts, an unreported feature neglected once across two different contexts may be more strongly inhibited than one neglected twice in the same context.

5. Modeling Accessibility and Precision

What is the nature of the observed changes in VLTM in Experiments 1-3? One possibility is that retrieval practice of one feature dimension enhances the accessibility of representations in that dimension relative to other unpracticed dimensions. Another possibility is that the observed memory effects can be explained by changes in the precision of representations in the practiced vs. unpracticed dimensions. To address this issue quantitatively, we estimated the relative contributions of decrements in accessibility and precision using a model-fitting method that partitioned errors in the final test phase into two component distributions reflecting: the probability of guessing (i.e., inversely related to accessibility) and the precision of retrieved feature representations.

5.1 Methods

On some proportion of trials during the test phase, response errors were very large. Such errors might reflect either a random guess or a very imprecise memory representation. The data can thus be described as a mixture of two distributions: a uniform distribution when the observer is unable to access the feature value and attempts to randomly guess from the full

space of feature values, and a von Mises (circular normal) distribution when the observer could access the representation in memory and made a noisy response centered around the true value. Maximum likelihood estimation can be used to recover the likelihood that responses were drawn from each of these distributions by fitting parameters that characterize each distribution (e.g., Zhang & Luck, 2008). Two parameters were of primary interest: (1) the height of the uniform component of the mixture, or the guessing rate (g); and (2) the standard deviation of the von Mises component (sd), or the inverse of memory precision.

We note one important caveat about our application of this technique: mixture modeling requires assumptions about the underlying distributions that are difficult to validate with a small number of observations. Specifically, we obtained fewer observations per observer and condition (48 in Experiment 1, 30 in Experiments 2 and 3) than previous studies that have used this technique (e.g., 150 in Zhang & Luck, 2008). To help compensate for this, we used a maximum likelihood estimation method in which we repeatedly estimated the best-fitting parameters from random starting values, and then extracted the modal value from this distribution. This iterative procedure was used to ensure that estimates of g and sd were as robust as possible (Figure 6A).

5.2 Results

5.2.1 Experiment 1
The model fitting procedure suggested that guessing occurred on 30.9% of test trials for Constant (95% CI for g: [25.4%, 36.3%]), 50.0% of trials for Switch ([43.0%, 56.2%]), and 48.2% of trials for Passive ([40.5%, 55.6%]). Guessing was less common for Constant than Switch ($t(35) = 7.04$, $p < 0.001$, $d = 1.17$) and Passive ($t(35) = 6.36$, $p < 0.001$, $d = 1.06$), which did not differ from each other ($t(35) = 0.646$, $p = 0.523$, $d = 0.108$). The precision of non-guess responses was similar across conditions: 15.9° for Constant (95% CI for sd: [14.1°, 17.8°]), 17.9° for Switch ([15.3°, 20.6°]), and 17.7° for Passive ([15.6°, 19.7°]). Indeed, Constant did not differ from Passive ($t(35) = 1.34$, $p = 0.189$, $d = 0.224$) or Switch ($t(35) = 1.24$, $p = 0.223$, $d = 0.207$), and they did not differ from each other ($t(35) = 0.152$, $p = 0.880$, $d = 0.025$). These results suggest that differences in the probability of successful access, rather than the precision of stored representations, may partly underlie the observed changes in VLTM performance due to prior retrieval practice. Individual differences in test phase performance provide further support for this interpretation (Figure 6B). RMSE was highly correlated with individual estimates of g (Constant: $r = 0.92$, $p < 0.001$; Switch: $r = 0.90$, $p < 0.001$; Passive: $r = 0.96$, $p < 0.001$), but not sd (Constant: $r = 0.15$, $p = 0.365$; Switch: $r = -0.13$, $p = 0.453$; Passive: $r = 0.006$, $p = 0.972$).

5.2.2 Experiment 2
The modeling results from Experiment 2 followed a similar pattern to Experiment 1. Guessing occurred on 31.3% of test trials for Constant (95% CI for g: [24.8%, 37.8%]), 44.4% of trials for Switch ([37.7%, 51.5%]), and 37.5% of trials for Passive ([29.6%, 45.4%]). Guessing was less common for Constant than Switch ($t(29) = 4.68$, $p < 0.001$, $d = 0.854$), with a trend for Constant less than Passive ($t(29) = 1.84$, $p = 0.077$, $d = 0.335$). There was also a trend for more guessing on Switch than Passive ($t(29) = 1.86$, $p = 0.073$, $d = 0.339$). The precision of non-guess responses was similar across conditions: 15.3° for Constant (95% CI for sd: [12.7°, 17.9°]), 17.80° for Switch ([13.4°, 22.2°]), and 16.3° for Passive ([14.1°, 18.66°]). No differences between conditions reached significance (ps > 0.335, ds < 0.179). As in Experiment 1, RMSE was highly correlated with individual estimates of g (Constant: $r = 0.85$, $p < 0.001$; Switch: $r = 0.88$, $p < 0.001$; Passive: $r = 0.92$, $p < 0.001$), but generally not with sd (Constant: $r = -0.14$, $p = 0.450$; Passive: $r = 0.16$, $p = 0.400$; Switch: $r = 0.26$, $p = 0.157$).

5.2.3 Experiment 3—Guessing occurred on 34.8% of test trials for Constant (95% CI for *g*: [27.2%, 42.4%]), 44.9% of trials for Switch ([38.0%, 51.8%]), and 50.6% of trials for Variable ([41.7%, 59.5%]). Guessing was less common for Constant than Switch (*t*(29) = 2.66, *p* = 0.01, *d* = 0.487) and Variable (*t*(29) = 3.49, *p* = 0.002, *d* = 0.636), which did not differ from each other (*t*(29) = 1.49, *p* = 0.146, *d* = 0.273). The precision of non-guess responses was similar across conditions: 22.5° for Constant (95% CI for *sd*: [16.0°, 29.1°]), 22.4° for Switch ([17.4°, 27.6°]), and 23.1° for Variable ([16.1°, 30.2°]). No differences between conditions reached significance (*p*s > 0.871, *d*s < 0.030). RMSE was again highly correlated with individual estimates of *g* (Constant: *r* = 0.88, *p* < 0.001; Switch: *r* = 0.82, *p* < 0.001; Variable: *r* = 0.93, *p* < 0.001). Conventional analyses suggested a significant, positive relationship between RMSE and *sd* for Constant (*r* = 0.41, *p* = 0.024), though not Switch (*r* = 0.27, *p* = 0.152) or Variable (*r* = -0.08, *p* = 0.654). However, inspection of the distribution of *sd* parameter fits revealed the presence of a single high-leverage observer in the Constant condition (*sd* = 96.3°, RMSE = 101.8°), whose *sd* estimate exceeded 3 standard deviations plus the mean *sd*. Without this observer, the correlation between RMSE and *sd* for Constant (*r* = 0.03, *p* = 0.872) was no longer significant. This result suggests that the apparent relationship between RMSE and *sd* in for Constant was spurious, likely due to poor parameter estimation in these small samples.

5.3 Discussion

The guessing parameter results from all three experiments suggest that retrieving one feature of an object from VSTM affects which features of that object are accessible in VLTM. At the same time, the precision parameter results suggest that VSTM retrieval does not affect the fidelity with which the features of an object are stored in VLTM. Strong correlations between RMSE and guessing but not precision in each condition further suggest that the main RMSE results from all experiments reflect differences in accessibility.

6. General Discussion

This study aimed to elucidate how internal attention to features in VSTM influences VLTM for novel objects. Specifically, we explored how retrieval history determines the weighting of features within individual objects over the course of learning. To address this question, we designed a visual memory task that manipulated feature retrieval history on an item-level basis. Observers initially viewed a series of shapes defined by unique values along three feature dimensions (location, orientation, and color), and reported one of these features from VSTM using a continuous report procedure. In a subsequent test phase, observers reported either the same feature or an unpracticed feature from VLTM. Across three experiments with different groups of observers, we discovered important consequences of internal attention for the organization of object memory.

In Experiment 1, we evaluated three hypotheses about how retrieving an object's features from VSTM would affect VLTM for that object: (1) reporting one feature might affect only that feature (practice-benefit hypothesis); (2) internal attention to one feature might benefit all features (object-based hypothesis); and (3) internal attention might induce competitive dynamics among features (feature-competition hypothesis). Consistent with the third hypothesis, for a given object, memory was both enhanced for practiced features and suppressed for concurrently unpracticed features.

Retrieval history may thus play an important role in how objects are construed in visual memory, extending views that emphasize only the stimulus properties of the current display (Alvarez & Cavanagh, 2004; Bays & Husain, 2008; Luck & Vogel, 1997; Olson & Jiang, 2002; Wheeler & Treisman, 2002). Moreover, previous work that has manipulated internal attention to items in short-term memory (Lepsien & Nobre, 2006; Nobre et al., 2004) has

focused on immediate consequences. In contrast, we tested the long-term consequences of such orienting for subsequent memory. We found that long-term memory biases for specific items can be acquired rapidly (i.e., after 2-3 exposures) and over a relatively large corpus (i.e., 90-144 novel objects, depending on the experiment). Moreover, although our findings seem to contradict a standard object-based view—that processing one feature of an object benefits processing of its other features—they may be consistent in another sense. Specifically, as expanded below, we interpret our findings as suggesting that all features of an object come to mind during VSTM retrieval in an object-based manner, and that which features are selected vs. inhibited determines how an object is later represented in VLTM.

In Experiment 2, we placed greater pressure on the hypothesized role of memory retrieval in shaping the organization of long-term memory for objects by equating the amount of sensory exposure to all feature dimensions, and controlling for motor response learning. While in Experiment 1 the memory probe did not contain any of the three critical features, the memory probe in Experiments 2 and 3 was identical to the original object except for the feature being probed. As a consequence, observers received concurrent veridical access to unpracticed dimensions during the practice phase. Despite greater exposure to unpracticed features, memory for these features remained worse when other features had been practiced (and exposed less). These findings show that the suppression in VLTM is a robust consequence when internal attention neglects features in VSTM, even when these features are externally available.

Intriguingly, the size of the Switch vs. Passive difference at test roughly doubled in magnitude in Experiment 2 compared to Experiment 1. One explanation is that the unpracticed features were stronger competitors in Experiment 2 because they were externally available in the probe rather than only internally active in VSTM. This heightened competition may have led to even further weakening of the unpracticed features when the practiced features were successfully reported. Indeed, the unpracticed features of the probe may have automatically captured attention simply because they matched the contents of working memory (Soto, Hodsoll, Rotshtein, & Humphreys, 2008). Just as external attention is recruited to filter out distractors, the perceptual availability of task-irrelevant features may have 'propped up' the unpracticed features in VSTM, requiring stronger inhibition via internal attention.

In Experiment 3, we aimed to probe how internal attention to multiple features affects learning about novel objects. Attending to two features in VSTM produced even greater suppression of the unpracticed feature than attending to one feature twice. What might account for this finding? We propose that memory retrieval functions as a form of reinforcement, and that competing representations that are repeatedly passed over for retrieval reinforcement are particularly susceptible to forgetting. Upon initial exposure to an object, there is no reason to privilege one feature dimension over another when forming a representation (absent systematic differences in salience or goals). Cued retrieval may disrupt this equal weighting by inducing competition among features vying for retrieval. Ultimately, the competition is resolved by the successful selection of the cued feature and the inhibition of uncued features. As a result, the retrieved feature enjoys a boost to its market share in long-term memory. Other dimensions that came to mind at the time of retrieval suffer a decrement, consistent with mechanisms of competition-based suppression that have been proposed to account for memory enhancement and suppression in other paradigms (Anderson et al., 1994; Anderson & Spellman, 2005; Norman & Newman, 2007). Experiment 1 is consistent with this account, since practiced features were remembered better than unpracticed features.

According to this framework, the representational strength of target and distractor features modulates the degree to which retrieval-based shifts in memory occur, such that unpracticed features that compete strongly with the practiced feature suffer greater impairment than weakly competing features. If the same feature is practiced for a second time, we expect unpracticed features to compete less fiercely during this second retrieval event because of their weakened status, and thus are not suppressed to the same extent.

If a different feature is probed upon the second exposure to an object, an accurate report of this feature may be more difficult because it was previously inhibited. Moreover, other previously unpracticed features may strongly compete during this second retrieval attempt, as their strength may approximate that of the probed feature, in addition to the relatively strong previously probed feature. As a consequence of this additional round of competition, never probed features may experience even further marginalization. Consistent with this prediction, in Experiment 3 we observed worse test performance on unpracticed features of objects for which multiple features had been previously practiced, relative to those for which a single feature was practiced multiple times (Figure 7).

Overall enhancement and suppression could reflect two different types of changes in memory: retrieval practice of one feature could reduce the *precision* or the *accessibility* of the representations of other features. These two changes are not mutually exclusive, and their relative contribution can be estimated from continuous report data using a standard mixture modeling procedure (Zhang & Luck, 2008). Each response is treated: (1) as arising from a representation with a certain precision when accessible, or (2) as a random guess when inaccessible. Modeling the data this way generates one parameter estimate representing the probability that responses were guesses (g) and a second parameter representing the standard deviation of the responses that were not guesses around the correct value (sd). This procedure was applied to the data from all three experiments. The parameter that most reliably differed between tests of practiced vs. unpracticed features was g, the probability of guessing. In contrast, the estimated precision of feature memory was not reliably different across conditions. These findings suggest that retrieval principally affects whether the features of an object are accessible, not the precision with which they are stored.

So far, we have interpreted enhancement and suppression in VLTM to be a consequence of selective feature retrieval from VSTM. An alternative possibility is that observers learned which feature was relevant for a given object during its initial presentation in the practice phase, and then strategically encoded this feature into VSTM on subsequent presentations, resulting in better VLTM performance in the test phase. According to this view, our findings may at least partly reflect differences in encoding rather than retrieval.

We consider this "selective encoding" account to be unlikely for the following reasons: First, observers were equally (Experiments 1 and 2) or more (Experiment 3) likely to be tested in VLTM on an unpracticed feature as they were on a practiced feature, reducing the incentive to strategically encode practiced features. Second, VSTM performance should have improved upon repeated presentations as observers learned over time which feature would be probed, but this improvement was not apparent after eliminating the possibility of motor learning in Experiments 2 and 3 (see Results). Third, VLTM performance should have been more related to VSTM performance on repeated (when selective encoding was possible) vs. initial presentations, but this difference did not emerge in any experiment (ps > 0.273). Fourth, even though observers were more likely to practice the same feature repeatedly in Experiment 3, VSTM performance was no worse when a new (Variable) vs. repeated (Constant/Switch) feature was probed in the second presentation ($p > 0.197$). Finally, insofar as the repeated practice of the same feature induced a selective encoding strategy, VSTM performance should have been better when the likelihood of repeated

practice was higher in Experiments 1 and 2 (100%) vs. Experiment 3 (67%), but no such differences emerged for the second presentation (ps > 0.361). Taken together, all of these findings are incompatible with the alternative selecting encoding interpretation of the observed VLTM effects.

Considering the role of memory strategies raises an interesting connection to the literature on directed forgetting (see MacLeod, 1998). For instance, after being presented with a VSTM array containing two objects, being post-cued to *forget* one of the objects leads to better memory for the remaining object in the array (Williams et al., 2013). Likewise, being post-cued to *remember* a subset of objects boosts VSTM for these objects relative to receiving no cue (Williams & Woodman, 2012). The relationship between instructing observers to maintain a feature in VSTM ('directed remembering') vs. to retrieve and report a feature from VSTM ('selective retrieval') awaits further study. These two modes of internal selection may have qualitatively similar consequences for memory (suppression and enhancement), but these consequences could arise from the same or different mechanisms. On the one hand, directed forgetting engages metacognitive strategies to voluntarily bias long-term memory, in contrast to the present case, where the primary task was to immediately report a feature from VSTM. On the other hand, similar executive control processes invoked to resolve competition may play a similar role in both kinds of enhancement/suppression, regardless of whether they are engaged intentionally (e.g., Anderson, 2003). Indeed, we have interpreted our findings in terms of competition among features during selective retrieval from VSTM, but the directed forgetting literature suggests additional mechanisms that should also be considered (e.g., MacLeod, 1998; Sahakyan et al., 2009).

The consequences of selective internal attention to features might extend beyond subsequent feature memory. For example, to the extent that retrieval history directly alters the weighting of features within object representations, perceptual sensitivity to previously retrieved features may be heightened upon subsequent encounters with that object (Goldstone, 1998). Relatedly, tuning of object representations may result in previously retrieved features appearing more salient, such that the presence of these features guides the deployment of attention across a crowded display to facilitate localization and identification. Finally, the competition induced in the current study may reflect the random and independent assignment of feature values to objects; introducing meaningful dependencies across dimensions may result in more generalized object-based benefits of feature retrieval. These are all consequences of selective feature retrieval for subsequent processing of the same object. Similar dynamics also operate at the level of objects themselves, with attention enhancing memory of selected objects and suppressing memory of inhibited objects (Fox, 1995; Lepsien & Nobre, 2006; Nobre et al., 2004; Tipper, 1995). An open question concerns how selective feature retrieval for one object biases processing of other objects, either when presented concurrently or encountered in the future.

7. Conclusions

Taken as a whole, the present results reveal the highly interactive nature of attention, learning, and memory, as well as how 'cognitive actions' can have long-lasting consequences for the organization of object representations. Although we focused on internal attention, cognitive operations that transform mental representations in the course of 'reading out' their contents appear to be pervasive. For example, when making a decision, although our preferences obviously guide our choices, our choices can also shift our preferences (Brehm, 1956; Festinger, 1957; Lieberman, Ochsner, Gilbert, & Schacter, 2001). Just as physical objects are neither static nor inert, objects of perception that comprise the internal visual world are amenable to cognitive manipulation. In this way, even

the simple act of selection during the initial stages of learning has the potential to leave enduring traces in memory.

Acknowledgments

References

Al-Aidroos N, Said CP, Turk-Browne NB. Top-down attention switches coupling between low-level and high-level areas of human visual cortex. Proceedings of the National Academy of Sciences. 2012; 109(36):14675–14680.

Alvarez GA, Cavanagh P. The capacity of visual short-term memory is set both by visual information load and by number of objects. Psychological Science. 2004; 15(2):106–111. [PubMed: 14738517]

Anderson MC, Bjork RA, Bjork EL. Remembering can cause forgetting: Retrieval dynamics in long-term memory. Journal of Experimental Psychology: Learning, Memory, and Cognition. 1994; 20(5):1063–1087.

Anderson MC, Bjork EL, Bjork RA. Retrieval-induced forgetting: Evidence for a recall-specific mechanism. Psychonomic Bulletin & Review. 2000; 7(3):522–530. [PubMed: 11082860]

Anderson MC. Rethinking interference theory: Executive control and the mechanisms of forgetting. Journal of Memory and Language. 2003; 49(4):415–445.

Anderson MC, Spellman BA. On the status of inhibitory mechanisms in cognition: memory retrieval as a model case. Psychological Review. 2005; 102(1):68–100. [PubMed: 7878163]

Andrews-Hanna JR, Reidler JS, Huang C, Buckner RL. Evidence for the default network's role in spontaneous cognition. Journal of Neurophysiology. 2010; 104(1):322–335. [PubMed: 20463201]

Awh E, Jonides J. Overlapping mechanisms of attention and spatial working memory. Trends in Cognitive Sciences. 2001; 5(3):119–126. [PubMed: 11239812]

Bays PM, Husain M. Dynamic shifts of limited working memory resources in human vision. Science. 2008; 321:851–854. [PubMed: 18687968]

Brady TF, Konkle T, Alvarez GA. Compression in visual working memory: Using statistical regularities to form more efficient memory representations. Journal of Experimental Psychology: General. 2009; 138(4):487–502. [PubMed: 19883132]

Brady TF, Alvarez GA. Hierarchical encoding in visual working memory: ensemble statistics bias memory for individual items. Psychological Science. 2011; 22(3):384–392. [PubMed: 21296808]

Carpenter SK, Pashler H, Vul E. What types of learning are enhanced by a cued recall test? Psychonomic Bulletin & Review. 2006; 13(5):826–830. [PubMed: 17328380]

Carrasco M, Ling S, Read S. Attention alters appearance. Nature Neuroscience. 2004; 7(3):308–313.

Carrier M, Pashler H. The influence of retrieval on retention. Memory & Cognition. 2002; 20(6):633–642. [PubMed: 1435266]

Chelonis JJ, Calton JL, Hart JA, Schachtman TR. Attenuation of the renewal effect by extinction in multiple contexts. Learning and Motivation. 1999; 30(1):1–14.

Chun MM, Golomb JD, Turk-Browne NB. A taxonomy of external and internal attention. Annual Review of Psychology. 2011; 62(1):73–101.

Chun MM, Johnson MK. Memory: enduring traces of perceptual and reflective attention. Neuron. 2011; 72(4):520–535. [PubMed: 22099456]

Chun MM, Turk-Browne NB. Interactions between attention and memory. Current Opinion in Neurobiology. 2007; 17(2):177–184. [PubMed: 17379501]

Ciranni MA, Shimamura AP. Retrieval-induced forgetting in episodic memory. Journal of Experimental Psychology: Learning, Memory, and Cognition. 1999; 25(6):1403–1414.

Crutcher RJ, Healy AF. Cognitive operations and the generation effect. Journal of Experimental Psychology: Learning, Memory, and Cognition. 1989; 15(4):669–675.

Curby KM, Glazek K, Gauthier I. A visual short-term memory advantage for objects of expertise. Journal of Experimental Psychology: Human Perception and Performance. 2009; 35(1):94–107. [PubMed: 19170473]

Fougnie D, Asplund CL, Marois R. What are the units of storage in visual working memory? Journal of Vision. 2010; 10(12):27. [PubMed: 21047759]

Fox E. Negative priming from ignored distractors in visual selection: A review. Psychonomic Bulletin & Review. 1995; 2(2):145–173.

Frigge M, Hoaglin DC, Iglewicz B. Some implementations of the boxplot. American Statistician. 1989; 43(1):50–54.

Goldstone RL. Perceptual learning. Annual Review of Psychology. 1998; 49(1):585–612.

Houghton, G.; Tipper, SP. A model of inhibitory mechanisms in selective attention. In: Dagenbach, D.; Carr, TH., editors. Inhibitory Processes in Attention, Memory, and Language. San Diego, CA: Academic Press; 1994. p. 53-112.

Johnson MK, Hirst W. MEM: Memory subsystems as processes. Theories of Memory. 1993; 1:241–286.

Johnson MK, Reeder JA, Raye CL, Mitchell KJ. Second thoughts versus second looks: An age-related deficit in reflectively refreshing just-active information. Psychological Science. 2002; 13:63–66.

Karpicke J, Roediger H. The critical importance of retrieval for learning. Science. 2008; 319:966–968. [PubMed: 18276894]

Klinger E, Cox WM. Dimensions of thought flow in everyday life. Imagination, Cognition and Personality. 1987; 7(2):105–128.

Lavie N, Fox E. The role of perceptual load in negative priming. Journal of Experimental Psychology: Human Perception and Performance. 2000; 26(3):1038–1052. [PubMed: 10884008]

Lepsien J, Nobre AC. Attentional modulation of object representations in working memory. Cerebral Cortex. 2006; 17(9):2072–2083. [PubMed: 17099066]

Levy BJ, Anderson MC. Inhibitory processes and the control of memory retrieval. Trends in Cognitive Sciences. 2002; 6(7):299–305. [PubMed: 12110363]

Luck SJ, Vogel EK. The capacity of visual working memory for features and conjunctions. Nature. 1997; 390:279–280. [PubMed: 9384378]

MacLeod, CM. Directed forgetting. In: Golding, JM.; MacLeod, CM., editors. Intentional forgetting: Interdisciplinary approaches. Mahwah, NJ: Lawrence Erlbaum Associates; 1998.

Mason MF, Norton MI, Van Horn JD, Wegner DM, Grafton ST, Macrae CN. Wandering minds: The default network and stimulus-independent thought. Science. 2007; 315:393–395. [PubMed: 17234951]

Melton AW. The situation with respect to the spacing of repetitions and memory. Journal of Verbal Learning and Verbal Behavior. 1970; 9:596–606.

Most SB, Scholl BJ, Clifford ER, Simons DJ. What you see is what you set: sustained inattentional blindness and the capture of awareness. Psychological Review. 2005; 112(1):217–242. [PubMed: 15631594]

Neill WT, Valdes LA. Persistence of negative priming: Steady state or decay? Journal of Experimental Psychology: Learning, Memory, and Cognition. 1992; 18(3):565–576.

Nobre AC, Coull J, Maquet P, Frith C, Vandenberghe R, Mesulam M. Orienting attention to locations in perceptual versus mental representations. Journal of Cognitive Neuroscience. 2004; 16(3):363–373. [PubMed: 15072672]

Norman K, Newman E. A neural network model of retrieval-induced forgetting. Psychological Review. 2007; 114(4):887–953. [PubMed: 17907868]

Nungester RJ, Duchastel PC. Testing versus review: Effects on retention. Journal of Educational Psychology. 1982; 74(1):18–22.

Olson IR, Jiang Y. Is visual short-term memory object based? Rejection of the "strong-object" hypothesis. Attention, Perception, & Psychophysics. 2002; 64(7):1055–1067.

Posner MI, Snyder CR, Davidson BJ. Attention and the detection of signals. Journal of Experimental Psychology: General. 1980; 109(2):160–174.

Rinaman, W.; Heil, C.; Strauss, M.; Mascagni, M.; Sousa, M. CRC Standard Mathematical Tables and Formulae. 30. CRC Press; 2012. Probability and Statistics.

Roediger HL, Karpicke JD. The power of testing memory: Basic research and implications for educational practice. Perspectives on Psychological Science. 2006; 1(3):181–210.

Sahakyan L, Waldum ER, Benjamin AS, Bickett SP. Where is the forgetting with list-method directed forgetting in recognition? Memory & Cognition. 2009; 37(4):464–476. [PubMed: 19460953]

Singer, JL. Daydreaming: An introduction to the experimental study of inner experience. New York, NY: Crown; 1966.

Slamecka NJ, Graf P. The generation effect: Delineation of a phenomenon. Journal of Experimental Psychology: Human Learning and Memory. 1978; 4(6):592–604.

Soto D, Hodsoll J, Rotshtein P, Humphreys GW. Automatic guidance of attention from working memory. Trends in Cognitive Sciences. 2008; 12(9):342–348. [PubMed: 18693131]

Theeuwes J. Perceptual selectivity for color and form. Attention, Perception, & Psychophysics. 1992; 51(6):599–606.

Tipper SP. The negative priming effect: Inhibitory priming by ignored objects. Quarterly Journal of Experimental Psychology. 1985; 37A(4):571–590. [PubMed: 4081101]

Treisman AM, Gelade G. A feature-integration theory of attention. Cognitive Psychology. 1980; 12(1): 97–136. [PubMed: 7351125]

Treisman AM. How the deployment of attention determines what we see. Visual Cognition. 2006; 14(4-8):411–443. [PubMed: 17387378]

Tukey, JW. Exploratory Data Analysis. Reading, MA: Addison-Wesley; 1977.

Uncapher MR, Hutchinson JB, Wagner AD. Dissociable effects of top-down and bottom-up attention during episodic encoding. Journal of Neuroscience. 2011; 31(35):12613–12628. [PubMed: 21880922]

Wheeler ME, Treisman AM. Binding in short-term visual memory. Journal of Experimental Psychology: General. 2002; 131(1):48–64. [PubMed: 11900102]

Wilken P, Ma WJ. A detection theory account of change detection. Journal of Vision. 2004; 4(12): 1120–1135. [PubMed: 15669916]

Williams M, Woodman GF. Directed forgetting and directed remembering in visual working memory. Journal of Experimental Psychology: Learning, Memory, and Cognition. 2012; 38(5):1206–1220.

Williams M, Hong SW, Kang M, Carlisle NB, Woodman GF. The benefit of forgetting. Psychonomic Bulletin & Review. 2013; 20(2):348–355. [PubMed: 23208769]

Xiao L-Q, Zhang J-Y, Wang R, Klein SA, Levi DM, Yu C. Complete transfer of perceptual learning across retinal locations enabled by double training. Current Biology. 2008; 18(24):1922–1926. [PubMed: 19062277]

Zhang W, Luck SJ. Discrete fixed-resolution representations in visual working memory. Nature. 2008; 453:233–235. [PubMed: 18385672]

Highlights

- We examined consequences of internal attention for visual learning of novel objects

- Reporting a feature from short-term memory boosts its long-term memory representation

- Other unreported features of the same object are suppressed in long-term memory

- Suppression is amplified when multiple features are reported from short-term memory

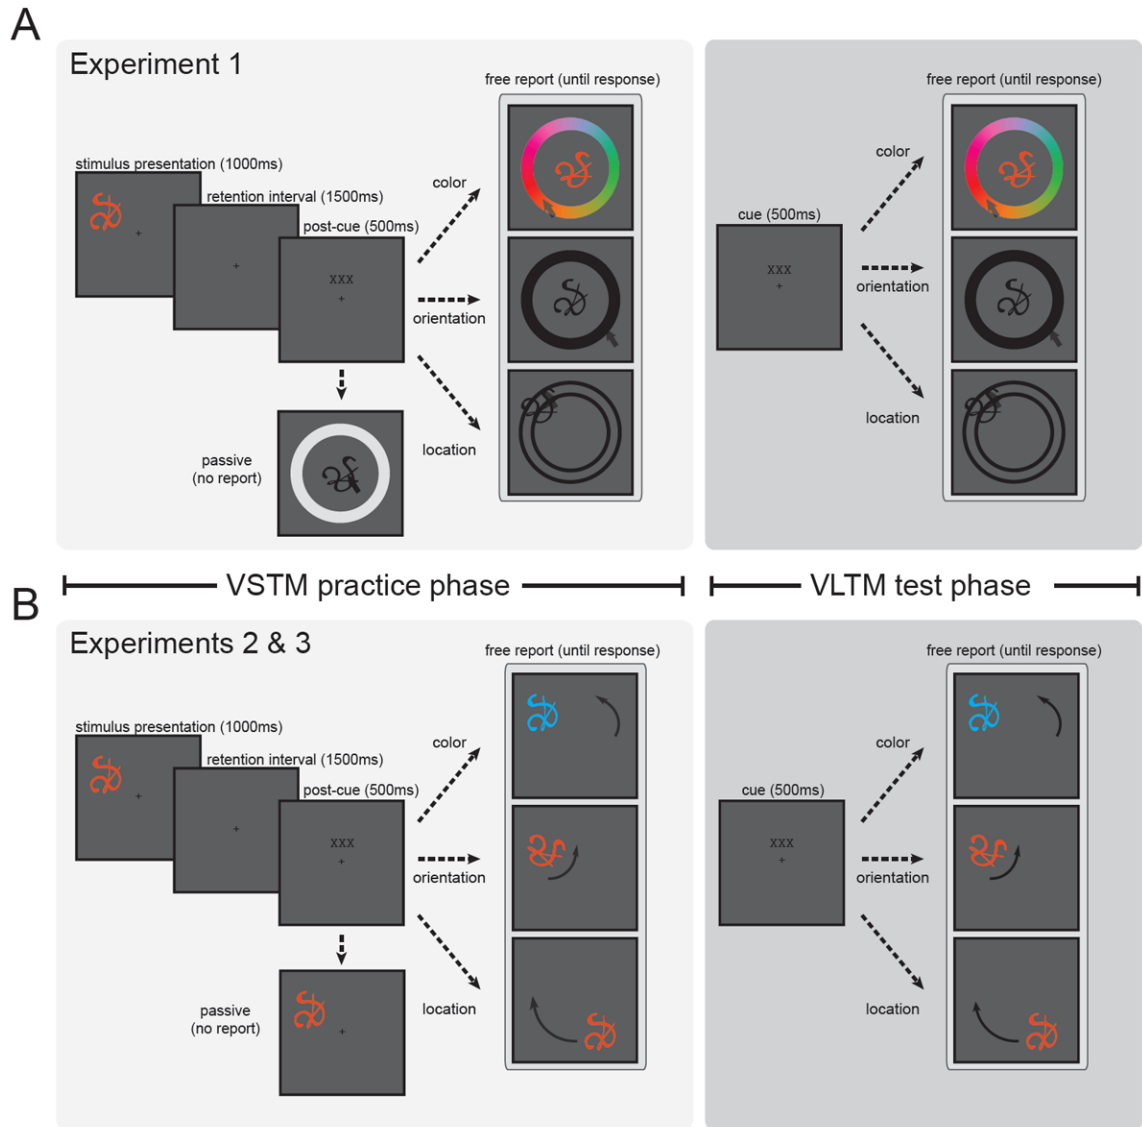- Internal attention modulates feature weighting in long-term object representations

Figure 1. Task display
On each trial of the practice phase, observers saw an object defined by a unique color, orientation, and location. After offset, observers were prompted by a post-cue ("xxx" replaced by "color", "orientation", "location", or "click") to report one of its features from VSTM. Observers reported this feature by continuously adjusting a memory probe until it matched the original stimulus. All objects were viewed multiple times during the practice phase. This was followed by a test phase in which observers could rely only on VLTM to report the features of a probe. The nature of the probe differed across experiments. **(A)** In Experiment 1, the memory probe was a black, canonically oriented version of an object presented in the center of the display. Observers used a visible response wheel to manipulate the memory probe along the cued dimension. **(B)** In Experiments 2 and 3, the memory probe was identical to the original object in appearance, except for the cued dimension. Observers adjusted the probe by moving the mouse until a good match was found. There was no response wheel, and the mapping between mouse position and feature space was perturbed from trial-to-trial.
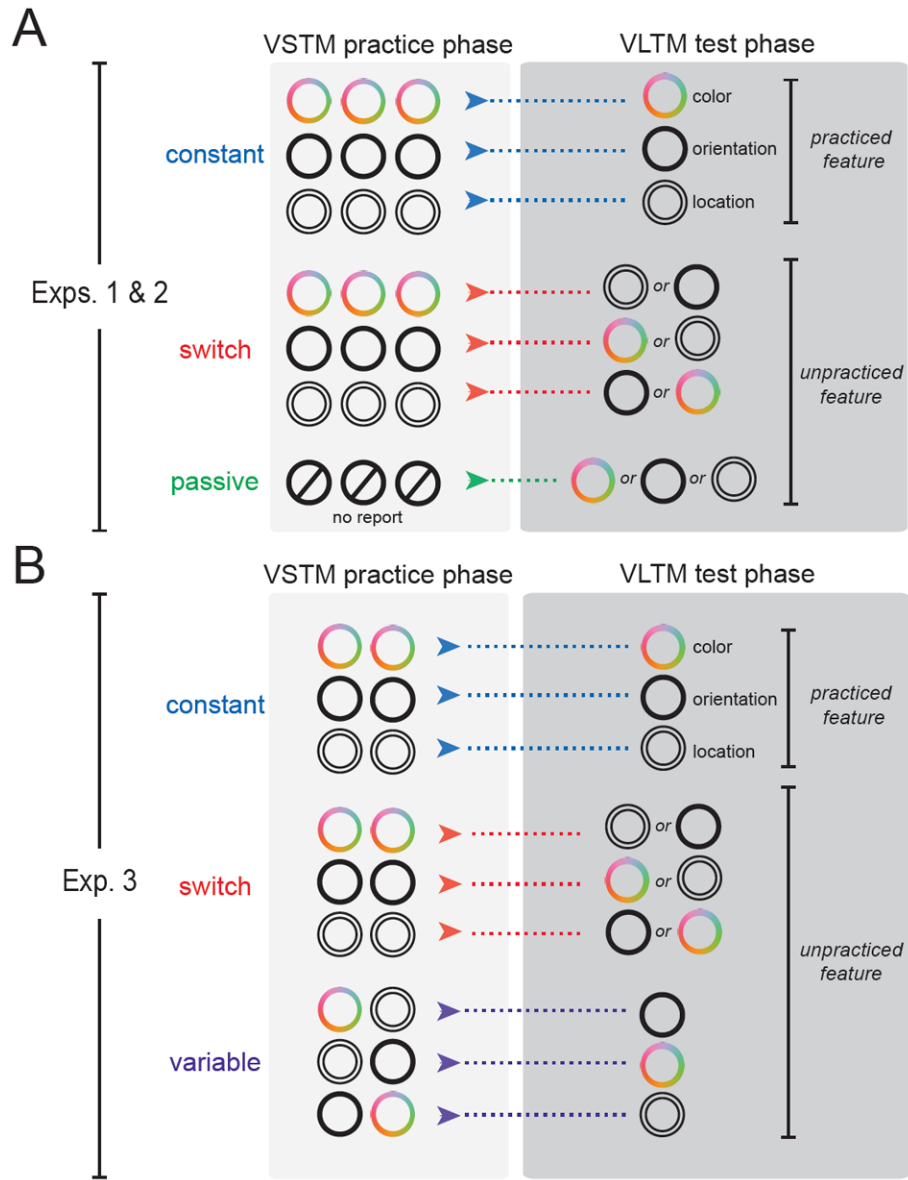
Figure 2. Design
(A) In Experiments 1 and 2, all objects were viewed three times during the practice phase. Each object was assigned to the Constant, Switch, or Passive condition. For Constant objects, the same feature was reported throughout both the practice and test phases. Switch objects were probed on the same feature in the practice phase but on an unpracticed feature in the test phase. Passive objects were viewed the same number of times during the practice phase, but no features were reported; one of the unpracticed features was subsequently tested. **(B)** In Experiment 3, all objects were viewed twice during the practice phase. Each object was assigned to the Constant, Switch, or Variable condition. Constant and Switch conditions were identical to before (but with two repetitions). For Variable objects, two different features were retrieved in the practice phase, and like the Switch condition, an unpracticed feature was subsequently tested.
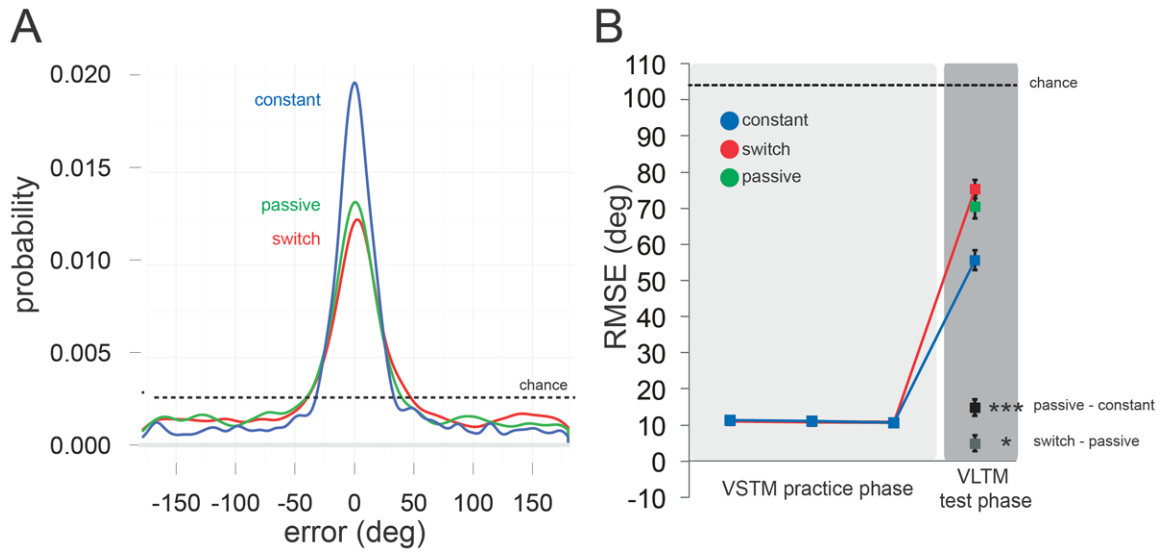
Figure 3. Results of Experiment 1
(A) Empirical distributions depicting the probability of observing an error of a given magnitude in the VLTM test for each condition. These distributions were constructed for visualization purposes only, by pooling data across all observers. **(B)** Task performance across VSTM practice and VLTM test phases as measured by RMSE for objects successfully reported during practice phase. In the test phase, practiced features (Constant) were reported most accurately, while unpracticed features of practiced objects (Switch) were impaired relative to unpracticed features of unpracticed objects (Passive). The height of the dotted line represents expected performance under random guessing. Error bars reflect ±1 SEM. *$p<0.05$, ***$p<0.001$.
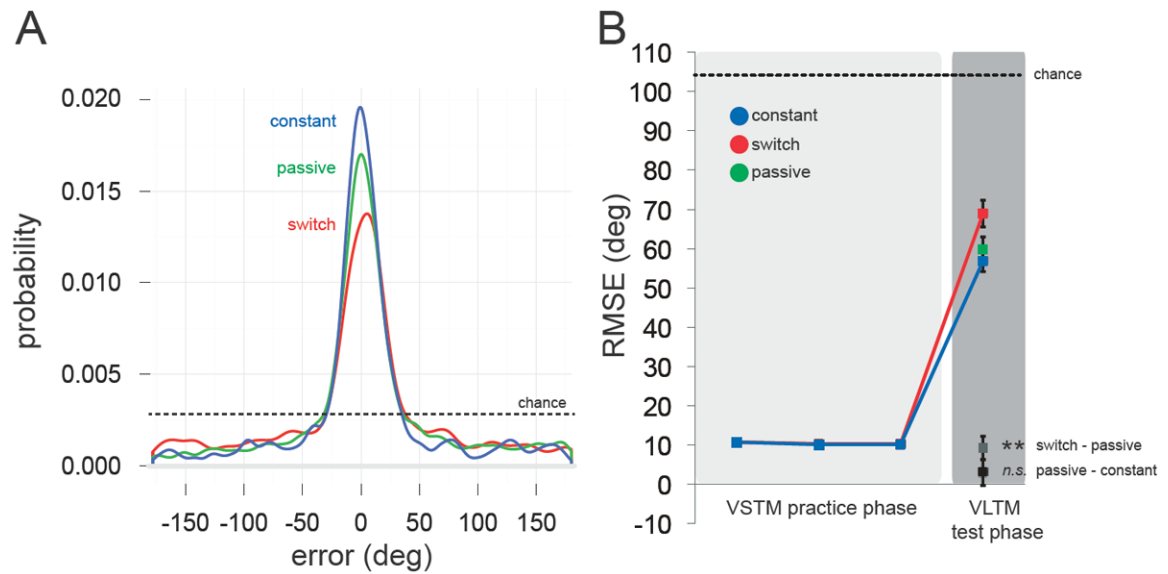
Figure 4. Results of Experiment 2
(A) Empirical distributions for each condition in the VLTM test phase, pooling data across observers for visualization purposes. **(B)** Task performance across VSTM practice and VLTM test phases as measured by RMSE. In the test phase, unpracticed features of practiced objects were impaired (Switch), even when observers received equal (vs. Passive) or greater (vs. Constant) veridical sensory exposure to these features during the practice phase. Error bars reflect ±1 SEM. **$p<0.01$.
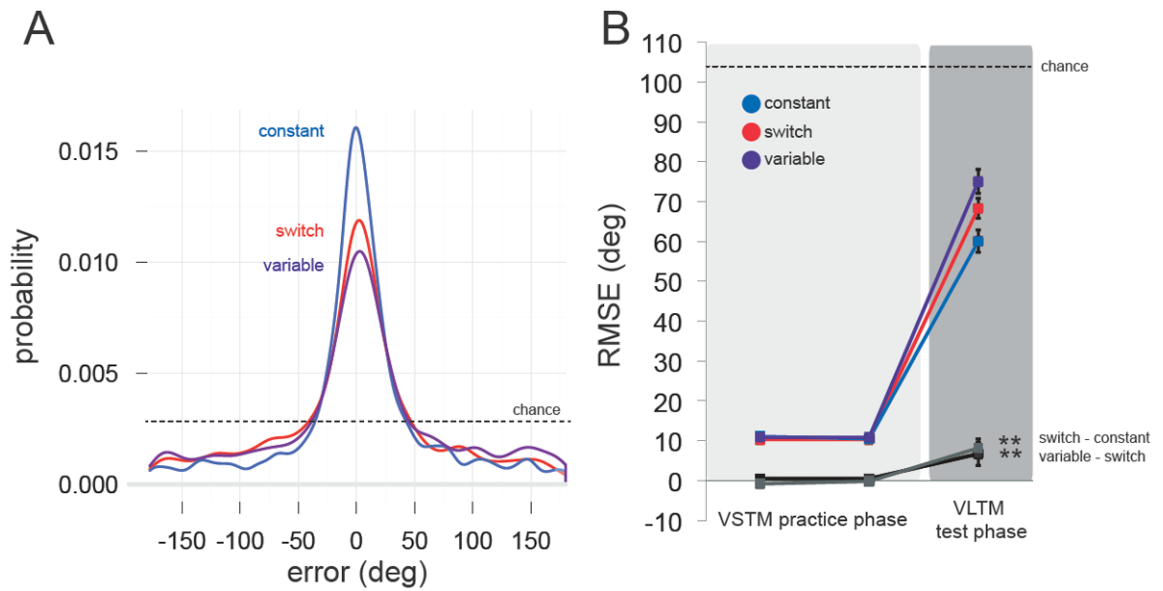
Figure 5. Results of Experiment 3
(A) Empirical distributions for each condition in the VLTM test phase, pooling data across observers for visualization purposes. **(B)** Task performance across VSTM practice and VLTM test phases as measured by RMSE. In the test phase, unpracticed features were impaired when belonging to an object for which two (Variable) vs. one (Switch) other features had been practiced. Error bars reflect ±1 SEM. **$p<0.01$.
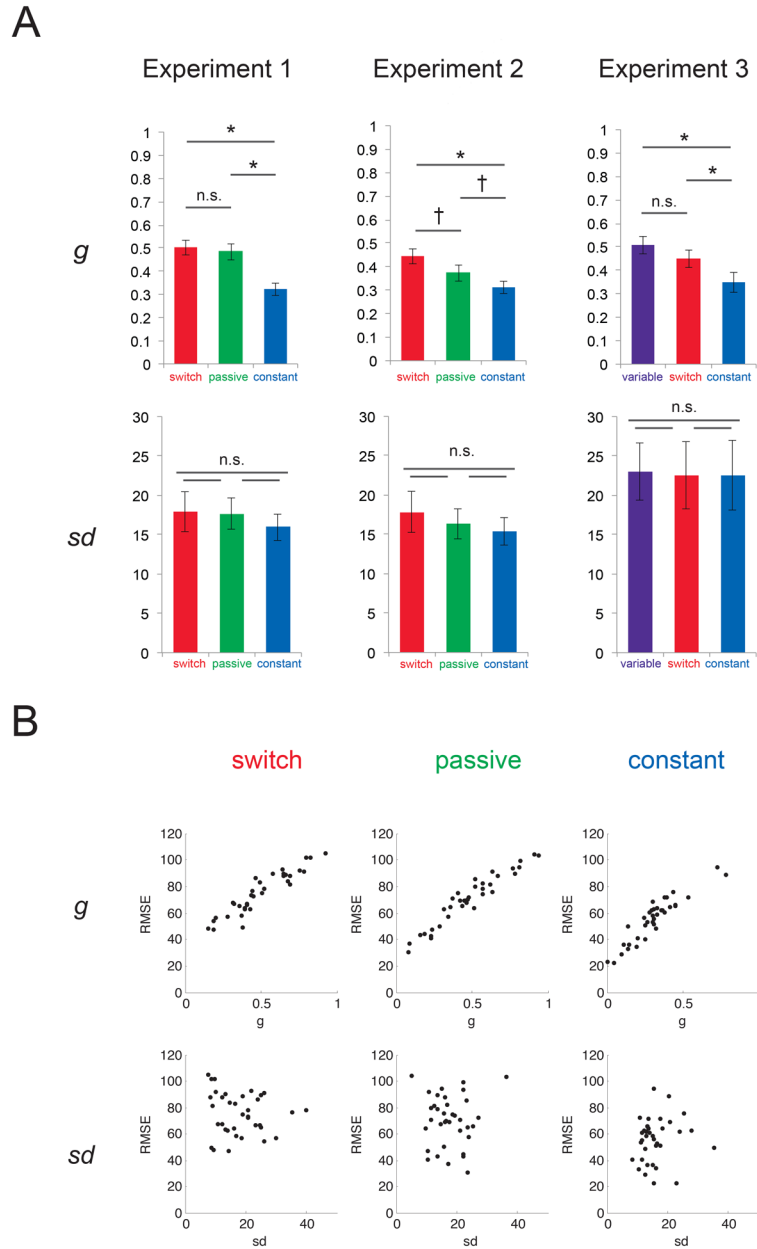
Figure 6. Model Fitting Results

(A) Guessing (*g*) and precision (1/*sd*) across conditions in the VLTM test phase of each experiment. The probability of guessing was generally reduced for practiced (Constant) vs. unpracticed feature (Switch/Variable/Passive) features, but the precision of non-guesses did not significantly differ between conditions. Error bars reflect ±1 SEM. * $p < 0.05$, † $p < 0.10$. **(B)** Relationship between parameter estimates and primary RMSE measure in Experiment 1. There were robust positive correlations between RMSE and *g*, but unreliable correlation between RMSE and *sd*. This pattern of correlations held across all experiments.
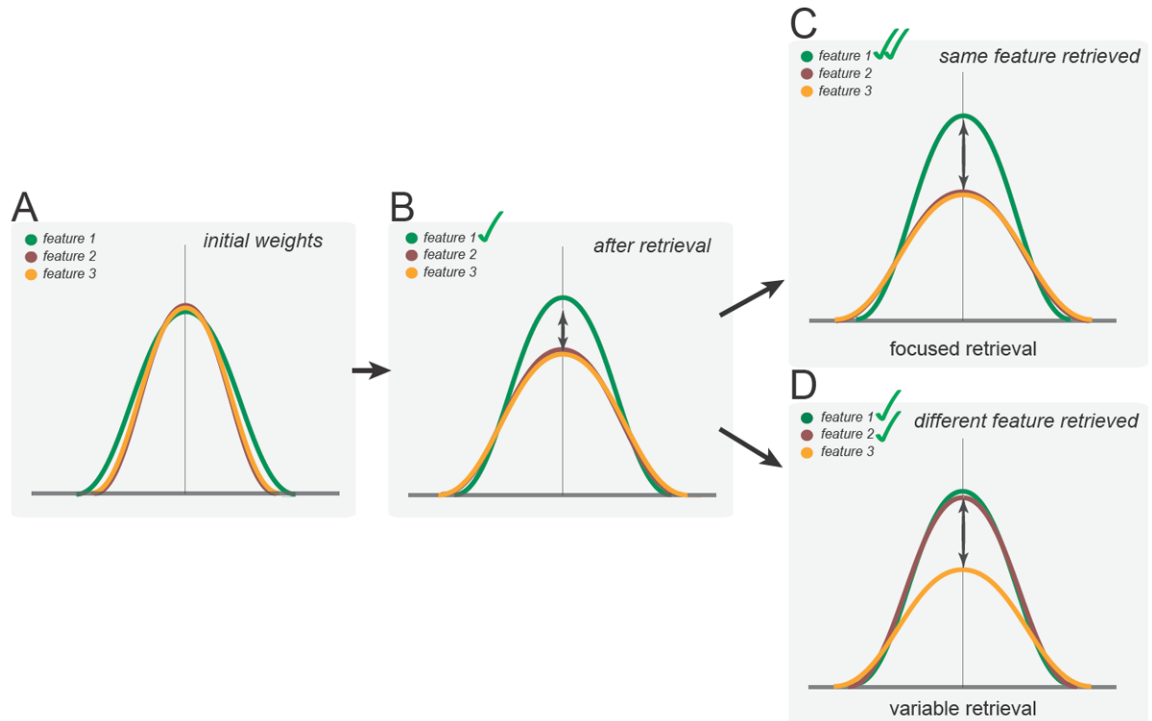
Figure 7. Interpretation
(A) Initially, attentional weights allocated to each feature dimension when forming a representation are roughly balanced. **(B)** Cued retrieval to a specific feature dimension may disrupt this equal weighting by inducing competition between dimensions for representation in short-term memory. This competition may be resolved via the selection of the object's value on the cued dimension and inhibition of values from uncued dimensions. Crucially, the consequences of this selective internal attention are not restricted to short-term memory, but extend to how various feature dimensions of an object are prioritized in long-term memory. **(C)** If the same feature is retrieved again from short-term memory in a subsequent encounter with that object, the competition among dimensions may be less fierce because this feature is already more strongly represented. As a result, the representation of the uncued dimensions may be further suppressed only slightly. **(D)** If a different feature is cued upon a subsequent encounter, accurate retrieval of this cued dimension from short-term memory may be more difficult due to the stronger status of the previously retrieved feature and prior inhibition of the newly cued feature. Moreover, given the weaker status of the cued feature, other previously uncued features may compete more, resulting in the further marginalization of still uncued feature dimensions in long-term memory.