Journal of Experimental Psychology: General

Real-World Objects Are Not Represented as Bound Units: Independent Forgetting of Different Object Details From Visual Memory

Timothy F. Brady, Talia Konkle, George A. Alvarez, and Aude Oliva

Real-World Objects Are Not Represented as Bound Units: Independent Forgetting of Different Object Details From Visual Memory

Timothy F. Brady
Harvard University

Talia Konkle
University of Trento

George A. Alvarez
Harvard University

Aude Oliva
Massachusetts Institute of Technology

Are real-world objects represented as bound units? Although a great deal of research has examined binding between the feature dimensions of simple shapes, little work has examined whether the featural properties of real-world objects are stored in a single unitary object representation. In a first experiment, we found that information about an object's color is forgotten more rapidly than the information about an object's state (e.g., open, closed), suggesting that observers do not forget objects as entirely bound units. In a second and third experiment, we examined whether state and exemplar information are forgotten separately or together. If these properties are forgotten separately, the probability of getting one feature correct should be independent of whether the other feature was correct. We found that after a short delay, observers frequently remember both state and exemplar information about the same objects, but after a longer delay, memory for the two properties becomes independent. This indicates that information about object state and exemplar are forgotten separately over time. We thus conclude that real-world objects are not represented in a single unitary representation in visual memory.

Keywords: binding, visual working memory, visual long-term memory, object-based attention

When we perceive a visual scene, we experience an organized and coherent set of objects and surfaces, not the disjointed patches of color or light that fall on the retina. We also appear to remember coherent, meaningful units: Moments after perceiving an office scene, for example, we might remember seeing objects such as a chair, a cup, and a person. In our subjective experience, it may seem that we perceive and remember each object as a coherent and integrated unit. However, a central question at the core of object representation is whether an object is actually represented as a completely bound unit, or whether it is represented with separable properties or dimensions.

Research on visual working memory has often claimed that the units of memory representation are bound objects (e.g., Cowan, 2001; Gajewski & Brockmole, 2006; Luck & Vogel, 1997; Vogel, Woodman, & Luck, 2001). For example, in their seminal study Luck and Vogel (1997) found that observers are equally good at remembering simple objects that vary along four features (color, size, orientation, and shape) as objects that vary along only a single feature (color or orientation alone), suggesting that working memory capacity may be limited by the number of objects rather than the number of visual features that can be stored. This suggests that the units of memory are bound object representations. Since Luck and Vogel, this strong object-based account of working memory representations has been shown to be too strong (e.g., Olson & Jiang, 2002; Wheeler & Treisman, 2002), but a significant amount of data nevertheless demonstrates a benefit to encoding multiple features of the same object (Fougnie, Asplund, & Marois, 2010; Luria & Vogel, 2011; Olson & Jiang, 2002; Xu, 2002), supporting the claim that visual working memory is at least partly object limited rather than limited only by storage of independent visual features (for a review, see Brady, Konkle, & Alvarez, 2011). However, there is also strong evidence that memory representations are not truly stored as a bound unit, and different features can be represented independently over short delays (Fougnie & Alvarez, 2011; Stefurak & Boynton, 1986) or long delays (Hanna & Remington, 1996).

Most of the work on the boundedness of object representation has examined only very simple objects made up of geometric shapes and colors. Much less work has examined whether real-world objects are represented as bound units. Because familiar real-world objects are more natural stimuli for the visual system, they might have more bound representations than objects that are made up of entirely dissociable low-level features that seem to be

stored independently even at the lowest levels of the visual system (e.g., orientation, color, spatial frequency; Magnussen, 2000) and that can be attended separately at encoding (Maunsell & Treue, 2006). Research on object recognition and long-term memory provide some proposals regarding the underlying representations of real-world objects (e.g., Diana, Yonelinas, & Ranganath, 2007; DiCarlo & Cox, 2007; Hummel, 2000; Riesenhuber & Poggio, 2000). In particular, these models typically assume "bound" representations of real-world objects. For example, view-based models of object representation tend to treat object representations as holistic, as storing a snapshot of an object from a particular view necessarily includes all the properties of that object in a single representation (Bülthoff & Edelman, 1992; Riesenhuber & Poggio, 1999; Tarr & Bülthoff, 1995). Further, most approaches to object recognition start from the assumption that object representations are independent of factors like lighting and size and rotation, but include all of the object's parts and properties together as the end-state of the ventral visual pathway (e.g., DiCarlo & Cox, 2007).

Similarly, long-term memory research typically makes a distinction between familiarity—a kind of holistic item memory—and recollection, or memory for the episodic details and context of an item (Diana et al., 2007; Ranganath, Yonelinas, et al., 2004). This distinction implicitly treats objects as unitary, where familiarity processes operate over object representations that do not require any binding, while recollection processes help retrieve information about how objects are bound to their contexts.

Thus, much of the existing literature—from object recognition, long-term memory, and visual working memory—treats real-world objects as though they are represented as a single bound unit. However, existing research on object representation does not directly address whether features of real-world objects are stored independently or as a single bound unit, and research from simple objects points to the possibility of independent representations of separate properties (Fougnie & Alvarez, 2011; Hanna & Remington, 1996; Stefurak & Boynton, 1986).

In the present study, we sought to empirically examine the assumption that different properties of a real-world object are represented as a single bound unit, as opposed to being represented as independent features. Furthermore, we sought to do so in a way that controlled for effects at encoding and retrieval that can falsely make objects look independent or bound (e.g., attention to only one property at encoding or being more attentive on some trials than others). We used the logic that, if a single bound representation of an object exists, then all of the object's features will be remembered or forgotten together. By contrast, if we observe that different properties are forgotten independently of each other, this would imply independent storage of the properties. We do not know the underlying features of object representation, but we can examine observers' ability to remember different object properties, such as an object's state, color, or exemplar (see Brady, Konkle, Alvarez, & Oliva, 2008). For example, observers can distinguish whether a door is open or closed (change of state), whether it is blue or red (change of color), or whether it is an ornate wooden door or a plain metal door (change of exemplar). Although these object properties are not likely to reflect primitive features for the visual system, they are semantically meaningful properties. We can thus use them to examine whether memory for different object properties may be supported in memory by different underlying features and thus whether different properties may be forgotten separately.

In Experiment 1, we varied object color and object state properties and examined whether observers forget one property more than the other over time. If so, this result would suggest that objects are not represented as single bound units with an all-or-none representation of object color and object state. In Experiment 2, we directly examined how memory for one property of an object (exemplar) was tied to memory for another property of an object (state) when observers were tested in a surprise memory task. In both experiments, which vary in paradigm, analysis method, and conditions of intentional and unintentional memory encoding, we found evidence for independent forgetting of different object properties. In Experiment 3, we showed that we can artificially induce apparently bound object representations by varying the strength of encoding for different objects, such that observers are likely to encode both properties of some objects and neither property of other objects. Together, these results demonstrate that real-world objects are not represented as a single bound unit in visual memory and that object representations can falsely appear bound after short delays because of encoding and retrieval factors.

Experiment 1

In a first experiment, we examined whether an object's state and an object's color are represented in a single bound representation or are represented as independent properties. To examine this, we looked at whether observers differentially forget these object properties over time. We had observers study real-world objects and then we tested their memory for the color and state of each object after either a short delay or a long delay. We reasoned that different rates of forgetting for different object properties would suggest that the properties were stored independently. For example, if the ability to detect both kinds of changes was equal at short delay but the ability to detect one kind of change decreased more than the other with increased delay, this would provide evidence that different object properties are forgotten separately.

Method

Participants. Forty-three naive observers were recruited from the Massachusetts Institute of Technology (MIT) participant pool (age range 18–35) and received $5 for their participation. All gave informed consent. Of the observers, 21 participated in the short-delay condition and 22 participated in the long-delay condition.

Stimuli. Object images were chosen from a previously published set of stimuli (Brady et al., 2008), supplemented with additional images from a commercially available database (Hemera Photo-Objects, Vol. I and II) and Internet searches using Google Image Search. Overall, 100 categorically distinct objects were selected, and for each of these objects two state-change images (differing in pose or configuration of parts) were selected. These objects were chosen such that they consisted of largely a single color, and this color was not intrinsic to the meaning of the object (e.g., the object would be recognizable in any color; see Figure 1 for example stimuli). In addition, we collected 200 other categorically distinct objects that differed from those in the main set but which were also recognizable in any color. These served as filler stimuli that would not be tested.

(a) Short-delay condition

Trial 1

Trial 2

Study display
1500ms
1000ms blank

Color test

State test

(b) Long-delay condition

Study phase

Test phase

1000ms
1500ms blank

Press key
for repeat

Color test

State test

Figure 1. Methods of Experiment 1. (a) In the short-delay condition, observers were shown three objects at a time and then tested on one of these objects. Either the color or the state of one of the objects was tested with a two-alternative forced choice, and which property of which object would be tested was not known in advance. (b) In the long-delay condition, all of the objects were shown, one at a time, and then observers were tested on one third of these objects. This test could be for either the color or the state of one of these studied objects. During the study period, observers also had to detect back-to-back repeats to ensure they were attending to the images. The two-alternative test displays were the same in the short- and long-delay conditions.

To create the final set of stimulus images, we rotated the hue of each object image to make it a random color. Hue is represented on a color wheel from 0 to 360 degrees, so the rotation required choosing a random angle for a given image and then adding that angle to the hue of each pixel for that image. Finally, for those images in the main set, we created two sets of images: one in the randomized color and one in a color 180 degrees in hue space from that color. Pilot testing showed that a rotation of 180 degrees in hue space led to an approximately equal degree of difficulty in the color forced-choice as in the exemplar forced-choice, and using such a large change in hue space also makes errors more likely to be due to cases of forgetting the object's color rather than to decreased memory precision (Alvarez, Konkle, Brady, Gill, & Oliva, 2009; Awh, Barton, & Vogel, 2007). This left us with a final set of 100 categories, each of which consisted of four images (two state-change images, each present in two different colors) and an additional set of 200 randomly colored filler objects.

Procedure.

General procedure. Observers were told to remember each object they saw as well as possible. Before the experiment began

they were given examples of the kind of forced-choice comparisons they would see, one example each of a state change and a color change. All observers sat approximately 60 cm from a 21-in. monitor. Objects were always shown at 7.5° visual angle. The experiment took approximately 20 minutes to complete.

Short-delay condition. Observers completed 100 trials each. Each trial started with a study display consisting of three objects arranged in a circle around a fixation cross. The objects were shown for 1.5 s, followed by a fixation cross for 1 s. Observers were then presented with a two-alternative forced choice. Two images were presented on the left and right side of the screen (see Figure 1), and observers were told to indicate which was present on the study display. Observers pressed Z if they had seen the left image and M if they had seen the right image. Then, after a brief (2.5 s) delay, the next trial began automatically.

The three items on the study display always contained two objects from the filler stimulus set and one object from the main stimulus set. The item on the subsequent test display was the one from the main stimulus set, presented in either two colors (color change condition) or two states (state change condition). The

location (left/right) of the correct answer was counterbalanced across observers, as was whether a given object was tested with a state change or a color change and which particular state and color image was the studied image and which was the foil image.

Long-delay condition. The long-delay condition consisted of a study phase and a test phase. In the study phase, observers were shown 340 objects, one at a time, for 1 s each at the center of the display. Of these objects, 100 were from the main stimulus set, 200 were filler objects, and 40 were back-to-back repeats of the filler objects. All items in the study stream were presented in a random order, and the subset of filler objects (40 of 200) that were repeated was randomized across observers. During the study phase, observers were told both to encode each object into memory and to press the space bar when an object repeated. This repeat-detection task served to ensure that observers were attending to each of the images.

In the test phase, observers completed 100 two-alternative forced-choice memory tests. Each of the 100 objects observers had seen from the main stimulus set was tested in the same manner as in the short-delay condition, again counterbalanced such that all objects were tested equally often in both the state and color conditions across observers.

Results

Two participants were excluded from the long-delay condition for failing to perform the back-to-back repeat detection task satisfactorily ($d' = 0.9$ and 2.5; mean d' for remaining participants was 4.2 with standard deviation 0.5). One participant was excluded from the short-delay condition for failing to complete the task. Thus, 20 participants from both the short- and long-delay conditions were entered into the final analyses.

Overall performance at the two-alternative forced choice task in the short-delay condition was 78.8% for color and 76.7% for state. In the long-delay condition, average performance was 67.8% for color and 72.5% for state. These results are plotted in Figure 2. Thus, there was a greater decrement in performance for color than for state with increasing delay (4% vs. 11%).

To examine the reliability of these effects, we ran a 2 × 2 mixed analysis of variance with state/color test as the within-subject factor and short/long delay as the between-subject factor. Participants had lower performance after a long delay relative to a short delay, main effect of delay, $F(1, 38) = 7.74$, $p = .008$, $\eta_p^2 = 0.14$. We found no main effect of test type, $F(1, 38) = 0.87$, $p = .35$, $\eta_p^2 = 0.02$, suggesting that neither state nor color tests were more difficult on average. However, we found a significant delay × test-type interaction, $F(1, 38) = 5.52$, $p = .02$, $\eta_p^2 = 0.13$, indicating a greater decrease in memory for color properties than for state properties as more time elapsed between study and test. In other words, color information is forgotten more than state information with a delay.

Discussion

We presented observers with real-world objects and then tested their memory for the objects' colors and state after either a short delay or a long delay. We found that observers' ability to detect a color change decreased markedly over time, whereas their ability to detect a change in object state remained relatively stable. These results demonstrate that observers do not forget each object as a bound unit but instead forget some properties more quickly than others. These findings suggest independent storage of different object properties in memory.

By including a short-term memory condition in addition to a long-term memory condition, we show that observers do not generally encode one feature preferentially over another (cf. Hanna & Remington, 1996) and that our test comparisons are equally difficult for observers in each dimension. Thus, we find that object color information is being lost from memory more quickly even though it is just as likely to be initially encoded and equally likely to be retrieved after a short delay.

An alternative interpretation of these results is that color information was never bound to object identity to begin with, even at a short delay. For example, in the short-delay condition, observers' may have simply remembered the three colors that were present on the study display without binding them to the identity of the

Figure 2. Results of Experiment 1. (a) Percent correct at forced-choice comparisons for state and color in both the short- and long-delay groups. After only a short delay, the color tests are slightly easier than the state tests. After a long delay, observers perform considerably worse on the color tests than the state tests. Error bars represent standard errors. (b) Decrement with delay for state and color. Observers' performance gets slightly worse for the state property with a delay but considerably worse for the color property.

objects. Such a lack of binding even in the short-delay condition could explain why color was lost more quickly than state information. Under this account, the different decay rates for color and state information occur because color was not bound to the object in the first place, whereas state information is more integrated with the object representation and is therefore retained longer. Importantly, this account is consistent with the main claim suggested here—namely, that objects are not stored in memory with all their features integrated into a unitary representation.

Our objects were purposefully chosen such that they did not have diagnostic colors (i.e., color was not a cue to the identity of the object). Thus, maintaining the color of each object in memory was expected to be difficult, as it was an arbitrary property that had to be bound into the object representation (and color is known to benefit object recognition when it is diagnostic but not when it is arbitrary; Naor-Raz, Tarr, & Kersten, 2003; Price & Humphreys, 1989; Tanaka & Bunosky, 1993). In the same way that objects that meaningfully connect to existing knowledge are easier to remember than objects that do not (e.g., Konkle, Brady, Alvarez, & Oliva, 2010; Wiseman & Neisser, 1964), meaningful features within an object may be easier to remember than arbitrary properties of objects. However, this manipulation does not necessitate our finding that different properties are forgotten at different rates. For example, we could have found that having to bind an arbitrary color into the object representation makes the entire bound representation more fragile or more likely to fall apart, resulting in a loss of performance in both the color and state comparisons. Alternatively, we could have found that it was difficult or impossible to match performance in the short-delay condition for the two dimensions. For example, observers may have had difficulty encoding arbitrary information such as color into memory in the first place. Instead, we find that the color information is initially encoded well but is then selectively forgotten while state information is preserved. This suggests that the actual underlying memory representation is not stored as a bound unit, even while controlling for independence resulting from encoding or retrieval factors.

Experiment 2

In Experiment 1, we used the fact that different properties of an object are forgotten at different rates to infer that the two properties are stored independently. In Experiment 2, we sought to examine more directly whether two object properties are remembered in a single bound representation or stored separately. To do so, we directly tested whether both properties of an object are remembered and forgotten together in long-term memory or whether they tend to be remembered and forgotten independently. For example, if observers remember seeing a glass of orange juice, do they systematically remember what kind of glass it was as well as how much juice was in it? How often do they forget only the shape of the glass or only the amount of juice? To examine this, we used two properties that have been shown to be forgotten at approximately the same rate (object state and object exemplar; see Brady et al., 2008) and looked at the probability of remembering one property given an observer remembered the other property about the same object (the *dependence* between the two properties).

In general, interpreting raw dependence scores (e.g., conditional probabilities) is complicated by a number of confounding factors

and is not a pure measure of how bound two properties are in the memory representation. For example, observers may be likely to either remember or forget both properties of an object because their overall attentiveness or fatigue level changed over the course of the experiment. Because both properties of the same object necessarily occur at the same point in time (they are both a part of the same object), this can make object representations look more bound than they truly are. In addition, successfully remembering one property may help in the retrieval of the other even if the properties are stored independently (e.g., encoding specificity; Tulving & Thomson, 1973), again introducing overestimates of boundedness.

On the other hand, the degree of boundedness can also be underestimated if there are differences in the difficulty of the exemplar and state comparisons. For example, any random variability in the degree of precision required for the state and exemplar comparisons across objects - causing state errors without exemplar errors for some objects, and vice versa for others - will masquerade as independent forgetting of features, underestimating the degree of boundedness in memory.

To avoid confounds from such encoding, retrieval, and stimulus factors, the critical manipulation in this experiment is to examine how memory performance changes over time. This holds confounding factors constant and also gives time for observers to forget some of the objects' properties. If the object properties are stored and forgotten independently, over time observers should be more likely to remember only a single property. In other words, the dependence between the object properties should decrease over time. In contrast, if the object properties are stored and forgotten together, then over time memory for the two object properties should have the same level of dependence. By taking into account how dependence changes with delay, we can observe not only how dependent the properties are on each other initially but also whether the objects are forgotten in a bound or an independent manner.

Thus, we tested observers' long-term memory performance after a short delay (30 min) and their long-term memory performance after a long delay (3 days) and examined whether the dependence between object properties decreased or stayed the same. Any decrease in dependence between the two object properties over time could not be caused by either encoding or retrieval factors, which were identical at the two delays. Thus, change in dependence over time allows us to infer whether two object properties are stored in a single unitary representation or are stored independently.

Method

Participants. Thirty naive observers were recruited from the MIT participant pool (age range 18–35) and received $5 for their participation. All participants gave informed consent. Of the observers, 15 participated in the short-delay condition and 15 participated in the long-delay condition.

Stimuli. Object images were chosen from previously published sets of stimuli (Brady et al., 2008; Konkle et al., 2010), supplemented with additional images from a commercially available database (Hemera Photo-Objects, Vol. I and II) and Internet searches using Google Image Search. Overall, 120 basic-level categories of object were selected, and for each of these categories

we selected two matching state images for each of two category exemplars. This yielded 120 object categories with 4 images each (2 exemplars × 2 states; see Figure 3).

Procedure.   The experiment consisted of a study phase and a test phase. In the study phase, observers were shown 120 objects one at a time for 200 ms each at the center of the display with an 1,800-ms interstimulus interval. During the presentation of the objects, they judged the physical size of the object (whether it was larger or smaller than a particular container they were shown, which was slightly smaller than a shoebox).

Following this task, they were given a surprise long-term memory task, either immediately following the study period (short delay) or after a 3-day delay (long delay). In the long-delay condition, observers were told immediately after the study period they would need to return in 3 days to perform memory tests. We used a surprise memory test and a 3-day delay to ensure that observers' performance was off ceiling at short delay and decreased substantially between the short and long delay, given that previous work has shown observers are quite good at these comparisons even after 5 hours of studying a large number of objects (Brady et al., 2008). To probe which properties of each object were encoded, we presented a four-alternative forced choice test display for each object, consisting of two exemplars (one familiar, one novel), each in two states (one familiar, one novel). Observers used the mouse to click on which of the four images they believed they had seen previously. After choosing an image, they separately reported how confident they felt (high or low) on both the state comparison and the exemplar comparison. The next trial then began automatically. There was no feedback.

Data analysis: Calculating the dependence score.   To address our main hypothesis, we examined the level of dependence between observers' reports of the state and exemplar properties. To do so, we calculated how much more likely observers were to get one property correct (e.g., state) if they got the other property correct than if they got it incorrect, taking into account the contributions of random guessing. In order to convert this into a dependence measure (% dependent), we first formalized two mod-

els: a fully independent model in which the properties are stored and forgotten independently and a fully bound model in which the properties are always stored and forgotten together. Then, we quantified where our observed data fell in between the predictions of the two models. Finally, for our critical comparison, we examined how this dependence score for the two properties changed between the short and long delays.

In the fully independent model (referred to as $D = 0$ below), there is never any benefit for memory of the state property given that exemplar was remembered, because the two properties are independent. Thus, no matter what the overall percent correct is, for an independent model of these two properties, the added memory benefit to one of remembering the other is 0:

$$P_{D=0}^{+}(state \mid exemplar) = 0$$

In the fully dependent model (referred to as $D = 1$ below), if the exemplar information is remembered, the state information will always be remembered. If all the objects are remembered, the increased memory performance for state information given exemplar information will go from chance (0.5) to remembered (1.0), for a maximal added benefit of 0.50. However, if observers do not remember an object, we assume they guess randomly from among the four items on the test display, and thus this guessing is independent for the two properties. As a consequence, even in the case of a fully bound underlying representation, random guessing for forgotten items will bring the added benefit down from 0.50. To account for this random guessing, we computed the guessing-adjusted fully bound model, based conceptually on that of Gajewski and Brockmole's (2006) model of boundedness in short-term memory, as follows.

First we estimate the percent remembered (R) for each property, based on the overall percent correct:

$$R(pc) = 2pc - 1$$

This formula treats memory as high threshold and takes into account that any overall percent correct (pc) was achieved not only

Figure 3.   Methods of Experiment 2. (a) Observers were presented with pictures of objects one a time. While they viewed the objects, their task was to indicate for each object whether it was smaller or larger than a container they were given. They were not told there would be a memory test for the objects. (b) In the short-delay condition, after finishing the size judgments for each object, observers were immediately told there would be a memory test and were tested on the objects they had seen after the study phase. In the long-delay condition, they were told there would be a memory test, but they came back in 3 days to complete it. Each test trial consisted of a four-alternative choice, with images of two different exemplars each in two different states. Observers' task was to click on which of the four images they had previously seen.

because items were remembered but also because items were sometimes forgotten but guessed correctly (Macmillan & Creelman, 2005). The "adjusted percent remembered" R estimates how often observers truly remember a property, after accounting for fortunate guesses, and is calculated based on overall performance and chance (here, 50% for each property).

For a given percent correct, the expected $p^+(state|exemplar)$ according to the bound model can then be calculated: Anytime observers remember the property ($R\%$ of the time), they should have complete dependence, $p^+(state|exemplar) = 0.5$, and anytime they forget a property ($1 - R\%$ of the time), guessing should cause complete independence, $p^+(state|exemplar) = 0$. Thus, although in theory a fully bound representation would have a $p^+(state|exemplar)$ of 0.5, once we take into account guessing, the dependence expected in a fully bound model (referred to as $D = 1$ below) varies as a function of overall percent correct (see Appendix A for derivation and simulation code):

$$p_{D=1}^+(state|exemplar) = \frac{R(pc)}{R(pc) + 1}$$

These expected dependences between properties in a fully independent model and in a fully bound model are plotted in Figure 4 as solid black lines.

Based on these models, for each observer we computed how dependent performance was between the state and exemplar conditions. This number could be a value between 0 (fully independent) and 1 (fully dependent), and it was computed based on the percentage of the way between the independent and bound model

predictions the observer's $p^+(state|exemplar)$ was at the observed percent correct. Because the fully independent model always predicts $p_{D=0}^+(state|exemplar) = 0$, this reduces to simply

$$D = \frac{p^+(state|exemplar)}{p_{D=1}^+(state|exemplar)}$$

where D is the dependence score of the observer, $p^+(state|exemplar)$ is how much more likely the observer was to get the state correct if he or she got the exemplar correct, and $p_{D=1}^+(state|exemplar)$ is the bound model prediction at the observer's percent correct.

Results

Two participants were excluded from the long-delay condition for failing to perform the size judgment cover task satisfactorily. Thus, 15 participants from the short-delay condition and 13 participants from the long-delay condition were entered into the final analysis.

Overall performance at the task was estimated separately for exemplar and state comparisons (e.g., whether observers indicated the correct exemplar independent of what state they chose and whether they indicated the correct state, independent of what exemplar they chose). In the short-delay condition, performance was 78.7% for exemplar and 72.3% for state. In the long-delay condition, average performance was 66.6% for exemplar and 63.0% for state.


Figure 4. Results of Experiment 2 for the state condition. (a) Given some level of performance in memory for the state of the object (*x*-axis), the *y*-axis shows how much more likely you are to remember the state if you remember the exemplar of the object than if you do not. If the two properties are completely bound in memory, the expected conditional probability is represented by the solid black line. If the two properties are remembered completely independently, the expected conditional probability would be 0 at every point on the *x*-axis. In the short-delay condition, the performance on state tests given the exemplar memory indicated some dependence, shown in the dashed red line. If memory for object state and exemplar information maintain this dependence relationship over a delay, performance in the long-delay condition would fall somewhere on this red dashed line. However, in the long-delay condition, there was a significantly lower dependence, plotted with the curve shown in dashed blue. Note that this figure shows the model fit to the group data for illustrative purposes; for purposes of analysis, the model was fit to each single participant's memory performance, and statistics were performed over the parameter estimates. (b) Observers in the short-delay condition have more dependence between the two properties than do observers in the long-delay condition, even after adjusting for the change in percent correct.

To address our main hypothesis, we examined the level of dependence between the two properties. After a short delay, we found that observers showed 46.6% dependence of state on exemplar (standard error of the mean [*SEM*] = 9.7%) and 27.4% dependence of exemplar on state (*SEM* = 4.9%), both significantly different from zero, $t(14) = 3.8$, $p = .0003$, Cohen's $d = 1.2$, and $t(14) = 5.6$, $p = .00006$, $d = 1.45$, respectively. The asymmetry between state|exemplar and exemplar|state is a result of the slightly different overall performance in the two conditions. After a 3-day delay, we found that observers showed 13.4% dependence of state on exemplar (*SEM* = 14.1%) and 7.6% dependence of exemplar on state (*SEM* = 8.7%), neither significantly different than zero, $t(12) = 0.97$, $p = .36$, $d = 0.27$, and $t(12) = 0.87$, $p = .40$, $d = 0.24$, respectively. In addition, these values of dependence were lower than those observed in the short-delay condition: state given exemplar, $t(26) = 1.99$, $p = .057$, $d = 0.75$; exemplar given state, $t(26) = 2.06$, $p = .05$, $d = 0.78$. These results show that the features were remembered more independently over time. Importantly, by comparing dependence rather than raw conditional probabilities, we remove the main effect of observers' decreased performance at a longer delay and adjust for the fact that guessing is necessarily independent.

Confidence. In addition to choosing which of the four stimuli they believed they had seen, observers gave us confidence judgments separately for the state and exemplar properties of the object. Thus, after observers chose their answer, we highlighted two of the objects (the one they chose and the change-of-state object) and they indicated how sure they were that the correct answer was the one they chose and not the other object (low or high confidence); then we did the same for the change-of-exemplar object. Overall, observers' confidence was well calibrated: Accuracy was higher when confidence was high (*M* = 84%, *SEM* = ±2.6% in the short-delay condition and *M* = 74% ±2.4% in the long-delay condition) than when confidence was low (61% ±1.7% and 58% ±1.3 in the short- and long-delay conditions, respectively).

Thus, we could also examine the degree of independence in memory representations by examining observers' self-reported memory strength rather than their percent correct. In particular, if observers said they had high confidence in one comparison (e.g., state), how likely were they to also have high confidence in the other (e.g., exemplar)? This metric is informative because it helps confirm that changes in guessing are not the source of the increased independence after a delay.

Participants' confidence decreased overall at long delay compared to short delay (chance of reporting high confidence: short delay, *M* = 64.9%, *SEM* = 4.6%; long delay, *M* = 46%, *SEM* = 4.2%). Importantly, this decrease appeared to be independent for state and exemplar properties. In particular, an observer's chance of reporting high confidence on one feature if he or she reported high confidence on the other feature was 82% (*SEM* = 3.1%) after a short delay and 62% (*SEM* = 3.9%) after a long delay. This decrease in dependence from short delay to long delay was significant, $t(26) = 4.04$, $p < .001$, $d = 1.53$. However, an observer's chance of reporting high confidence if he or she reported low confidence in the other feature did not differ across delay (short delay: *M* = 38.9%, *SEM* = 6.1%, long delay: *M* = 31.9%, *SEM* = 3.4%), difference not significant, $t(26) = 0.95$, $p = .35$, $d = 0.36$. To compare how much more likely participants were to report high

confidence on one feature given they did so on the other, we computed an odds ratio. In the short-delay condition, the odds ratio was 10.9 (±1.1), reflecting observers' having 10 times the odds of reporting high confidence in one property if they had high confidence in the other; in the long-delay condition this odds ratio was only 4.1 (±1.1), a significant difference ($p < .01$). This indicates that observers' likelihood of having high confidence for *both* features decreased with delay, even after discounting the general tendency for observers to have lower confidence after a delay.

Thus, these data show that observers' confidence also grows more independent with time, as does their percent correct. As delay time increases, high confidence on one feature is less likely to co-occur with high confidence on the other feature.

How independent are these features? The results indicate that state and exemplar information features are forgotten at least partially independently and that the underlying representation of these features is not fully bound. Can we quantify how independent the underlying memory representations for state and exemplar information must be from these data?

In the short-delay condition, we find a 46% dependence of state on exemplar. At one extreme, this could mean that memory representations for state and exemplar information are 46% overlapped; however, some percent of this dependence could also be attributed to "encoding correlation" (e.g., how likely observers are to encode or retrieve both properties of a given object due to other factors, such as attention, even though the underlying memory representations are actually independent). Thus, there is a spectrum of possibilities for the true feature independence, shown in Figure 5 (red line), with 46% boundedness and 0% encoding correlation at one extreme and 0% boundedness and 46% encoding correlation at the other (see formal model specification in Appendix B).

In the long-delay conditions, we find a 14% dependence of state on exemplar memory. The same logic also allows us to break down the 14% boundedness observed into components that reflect both

Figure 5. Best fit models that vary in both the memory boundedness and the correlation between state and exemplar encoding. Both the short-delay (red line) and long-delay (blue line) conditions can be fit by assuming a correlation at encoding, a bounded memory representation, or any mixture of the two. However, the combined data can be fit only by assuming that nearly all of the forgetting is independent (e.g., that the dependence results from correlations at encoding, indicated by the black X).

the true overlap in the features and the correlations due to encoding conditions, after taking into account the forgetting observed after the short delay. The full spectrum is shown in Figure 5 (blue line).

Critically, we can find the combination of feature boundedness and encoding correlations that simultaneously fit both the short and long delay conditions. The combined fit is shown in Figure 5 (black X). The only model that fits both delay conditions requires nearly 100% independent forgetting of the features.

The best fit parameters for explaining both the short delay and long-delay data suggest 97% independence between state and exemplar properties with $r = .63$ correlation of the likelihood of initially encoding the two properties about the same object. In fact, when simulating what we would expect from complete independent forgetting (100% all of the initial dependence coming from encoding), we find that our dependence measure should go from 46% to 14% when percent correct drops from 72% to 63%. This almost exactly matches the data we observe in Experiment 2. Thus, the current experiment provides strong evidence that the forgetting is almost totally independent, even though there is an initial dependence in how likely observers are to remember both the state and exemplar properties after a short delay.

Discussion

In Experiment 2, observers were required to remember both the state and exemplar of an object, and we examined whether observers remember both properties together or whether they forget the two properties separately. To examine this, we calculated the conditional probability of remembering one property given successful memory for the other, taking into account the independence of guessing. Interpreting a given level of dependence between the properties is difficult because this dependence can be influenced by a number of encoding and retrieval factors rather than simply the boundedness of the representation. Thus, we examined the change in this dependence over time, reasoning that any increase in independence of the two object properties with increased delay could not be caused by either encoding or retrieval factors, which were identical at the two delays, and must be caused by independent forgetting of the properties over time.

We found that at short delays, there was significant dependence between the two properties—observers were more likely than chance to remember both the state and exemplar of a particular object (46% bound). However, at a long delay this dependence is markedly decreased (14% bound). This suggests that much of the initial dependence observed at short delays is due to encoding or retrieval factors, such as the fact that all the features of a given object are presented at the same time and spatial location. However, ultimately the two properties of an object are forgotten separately. A straightforward model of how much of the dependence was caused by correlations at encoding versus boundedness in the memory representation indicated that almost all of the dependence was due to correlations at encoding, as forgetting appeared to be completely independent. This suggests that the initial dependence we observe between the two properties may be solely due to encoding and retrieval factors like differential attention on different trials. Thus, our data suggest almost totally independent storage of different object properties in memory.

The model of memory used in the calculation of dependence is a high-threshold model. We believe this is reasonable, even if it is

not entirely uncontroversial (e.g., Parks & Yonelinas, 2007; Wixted, 2007). In particular, we are asking people to recollect specific details of the objects, and our forced-choice comparisons require observers to choose between objects that are quite different from each other. Such large differences between the correct item and foil item seem to result in all-or-nothing memory retrieval in other features, like color (Brady, Konkle, Gill, Oliva, & Alvarez, 2012). However, our conclusions do not depend on the high-threshold nature of the model. In particular, forced-choice tasks in general are less sensitive to the distinction between signal detection and high-threshold models because with little bias, there is little dependence on the particular shape of the response operating characteristic curve that distinguishes these models (Macmillan & Creelman, 2005). Furthermore, in Appendix C we show that even if the underlying memory signal is better characterized by signal detection, our high-threshold model nevertheless does a reasonable job of characterizing the dependence between the properties. There seems to be no crucial difference between a high-threshold and signal detection model in examining the dependence between properties, which inherently depend not on the model of successful versus unsuccessful memory performance (the distinction between signal detection and high-threshold models) but on the dependence between performance with one stimulus and another.

Experiment 3

The data from Experiment 2 strongly suggest that the forgetting of separate features is independent and that the initial dependence we observe between properties at short delays is driven by encoding factors that lead observers to encode both properties about some objects and neither property about others. We hypothesized that this correlation in initial encoding probability could result from attentional differences over the course of the experiment. Thus, observers may be likely to either encode or fail to encode both properties of an object because their overall attentiveness or fatigue level changes over the course of the experiment. Because both properties of the same object necessarily occur at the same point in time and space, this could make object representations appear more bound than they truly are, particularly at short delays. In Experiment 3, we sought to test this directly by artificially increasing the likelihood of such encoding disparities. That is, we varied the display time of the objects, where some objects were presented for a longer duration than other objects. We expected the degree of dependence to be the same for short-presentation and long-presentation items analyzed separately but that combining the data across presentation durations would make object representations appear more dependent. Such a pattern would highlight the fact that variability in the quality with which different objects are encoded will artificially inflate the observed dependence between features—a limitation we addressed in Experiments 1 and 2 by measuring forgetting over time.

Method

Participants. Eleven naive observers were recruited from the Harvard University participant pool (age range 18–35) and received $5 for their participation. All participants gave informed consent. None of the participants had taken part in Experiment 2.

Procedure. The experiment was identical to the short-delay condition of Experiment 2, with one exception: For each observer,

a random half of the objects were displayed for a short duration (150 ms) and a random half of the objects were displayed for a long duration (500 ms). All other methods were identical to Experiment 2.

Results

As in Experiment 2, overall performance at the task was quite good, with performance at 85.2% correct for exemplar and 76.1% correct for state. The dependence scores computed separately for the short-presentation and long-presentation conditions were not significantly different, state|exemplar, 34% versus 45%, $t(10) = 0.62$, $p = .55$, $d = 0.19$; exemplar|state, 29% versus 31%, $t(10) = 0.14$, $p = .88$, $d = 0.05$, and on average were comparable to the short-delay condition of Experiment 2. However, when all the items were analyzed together, we found a 77.4% dependence of state on exemplar ($SEM = 8.5\%$) and a 43.6% dependence of exemplar on state ($SEM = 5.6\%$). These dependence levels were greater than the dependence scores observed in the short-encoding duration and long-encoding duration computed separately, state|exemplar, $t(10) = 4.3$, $p = .002$, $d = 1.29$; exemplar|state, $t(10) = 1.7$, $p = .11$, $d = 0.51$, as well as those observed in Experiment 2 (46.6% and 27%, respectively): state|exemplar, $t(24) = 2.3$, $p = .03$, $d = 0.90$; exemplar|state, $t(24) = 2.2$, $p = .04$, $d = 0.86$. These results suggest that increasing the disparity in how well particular objects are encoded can artificially increase the estimated dependence between properties.

Discussion

In Experiment 3, we manipulated whether observers had more or less time to study an object, in order to simulate the effects of stronger or weaker encoding that might naturally happen when studying a stream of items presented for equal durations. Critically, we found that feature representations were equally dependent whether items were presented for a short or long duration but that combining these trials together leads to an increased dependence estimate.

These data suggest that factors such as differential attention on different trials can cause dependence between properties. This highlights the fact that the observed dependence between two properties can be driven not only by the true underlying dependence but also by encoding and retrieval factors like differential attention on different trials. These results help put the results of Experiment 2 into context. In Experiment 2 we observed 46% dependence at the short delay but only 14% dependence at the long delay. We assume that the true underlying dependence of these two features does not change over time; thus, these two properties are at most 14% bound, with simulations putting this number closer to 0% bound. The results of Experiment 3 show how the 46% boundedness observed in short delay could be so high due solely to the contributions of encoding and retrieval factors. Together with those of Experiments 1 and 2, these results suggest almost totally independent storage of different object features in memory.

General Discussion

Across three experiments, we investigated whether different properties of real-world objects are represented with a single unitary object representation or whether they are represented independently. In Experiment 1, we showed observers arbitrarily colored real-world objects in different states and tested their memory for these properties immediately or after a delay. We found that, over time, arbitrary color information about the object was forgotten much more rapidly than the more meaningful state information. For example, people remembered they saw an upright lawn chair (as opposed to a reclined lawn chair) but not whether it was yellow or blue. This suggests that objects are not forgotten as bound units but instead that some object properties are forgotten more quickly than others.

In Experiment 2, we showed observers a set of categorically distinct objects that varied in two dimensions (object exemplar and state). We then probed observers' memory for state and exemplar information after either a short delay or long delay. After a short delay, observers frequently remember both properties about an object, but after a long delay memory for these properties was more independent. For example, after a short delay people were likely to remember that the cookie they saw was a chocolate chip cookie with a bite out of it; however, after more time, they were prone to confuse the cookie with an oatmeal-raisin cookie (exemplar information forgotten) but still remember that the cookie they saw had a bite out of it (state information remembered) or vice versa. This suggests that different object properties are forgotten independently over time, even within the same object. In fact, the forgetting we observed over time in Experiment 2 appeared to be almost entirely independent for the two properties.

In Experiment 3, we asked whether the initial dependence we observe between the properties at short delays derives from encoding and retrieval factors like attentional differences over the course of the experiment. We increased the heterogeneity of the initial encoding of the objects by showing some for shorter durations and some for longer durations. We found that increasing encoding disparity among objects leads to an inflated estimate of dependence.

Together, these data indicate that observers do not store a single unitary object representation in memory: Instead, some object properties persist while other properties are forgotten, and observers tend to forget different properties independently of each other for individual objects. Furthermore, although there is often a dependence between how likely observers are to remember different properties of the same object, we show that this is likely to be due to encoding factors rather than to reflect a bound underlying memory representation.

Below we discuss how independent storage of different object properties can have important repercussions for theories that advocate binding both in visual working memory and in long-term memory, as well as for models of object recognition, all of which tend to assume unitary object representations.

Object Representations Are Not Unitary

The existing literature on object binding in perception and visual working memory has tended to focus on perceptual binding, in particular, the binding of different low-level features such as orientation and color into coherent, bound objects. For example, feature integration theory proposes that we recognize low-level visual features such as color and orientation in parallel across the visual field, but binding these features together into coherent

objects requires attention (Treisman, 1998). Given the role of attention in this perceptual binding, it may not be surprising that much of the literature on visual working memory finds that features seem to be bound into objects in memory (e.g., Luck & Vogel, 1997), because (a) attention tends to be directed to all of the features of a particular object once that object is attended (O'Craven, Downing, & Kanwisher, 1999; Scholl, 2001) and (b) those objects we attend are likely to be the ones we remember (e.g., Chun, 2011; Rensink, O'Regan, & Clark, 1997). Thus, attention may be one encoding factor that often makes object representations appear bound, particularly in perception or after a short delay: If a particular object is attended, all of its features are attended, and those features will all be remembered well; by contrast, all the features of an unattended object will not be well remembered.[1] This role of attention during encoding could make even representations that are inherently independent and separable appear to be bound, as we found in Experiment 3. In addition to attention, other encoding and retrieval factors, such as the fact that successfully remembering one property may help in the retrieval of the other even if they are stored independently (e.g., encoding specificity; Tulving & Thomson, 1973), are all likely to impact the degree to which two independent properties of an object look bound.

Thus, we believe that existing evidence suggesting bound representations in visual working memory (e.g., Gajewski & Brockmole, 2006) may reflect, at least in part, shared encoding factors rather than truly unitary memory representations. In support of this idea, recent evidence suggests that observers may often remember one feature of an object but not another, even in simple stimuli like colored oriented lines (e.g., Bays, Wu, & Husain, 2011; Fougnie & Alvarez, 2011; Stefurak & Boynton, 1986). Further, remembering multiple features of the same object can come at a significant cost relative to remembering only a single feature (Fougnie et al., 2010).

In addition, it is possible to observe independence between features like color and orientation even in long-term memory. For example, observers can remember which shapes they saw without any impairment from a change in the color of the object between study and test (Hanna & Remington, 1996), suggesting independent representations of these features. However, because they used simple low-level features and told observers in advance what the memory tests would be like, Hanna and Remington may have caused observers to attend to only a single property of the visual objects during the study phase (e.g., using feature-based attention: Maunsell & Treue, 2006). Thus, their results could reflect encoding strategies rather than independence in the underlying memory representations. Similarly, encoding the features independently could play a role in the independence observed in much of the existing work in visual working memory (e.g., Stefurak & Boynton, 1986). In the present experiments, we ensured that both properties were equally relevant to the observer and still found independent forgetting of these properties.

By examining binding with real-world objects we were able to examine memory not only in the short term but also at longer intervals. This is challenging to do when using meaningless or simplified stimuli, as in previous approaches. In addition, by not telling people in advance about the memory test (in Experiments 2 and 3) and using properties that are not low level and thus cannot be attended too separately, we can avoid the potential for observers

to selectively encode one property over another. Thus, we believe that the method used in the current experiments—not only examining dependence between features at a single delay interval but also examining how it changes over time—may be critical to understanding whether seemingly bound representations are just a consequence of encoding and retrieval factors rather than a reflection of the true underlying structure of memory. By examining a change in dependence between properties over time, this approach allows us to examine the structure of memory representations while holding constant any dependence between properties induced by encoding and retrieval factors.

Binding and Perceptual Integrality

In the present experiments, we use object properties that are relatively high level: object state, object exemplar, and object color. This allows us to examine whether our memory representations for different properties are unitary, rather than whether our perception of two properties is unitary, as in the classic distinction between integral and separable dimensions (Garner, 1974). For example, using simple stimuli it can be shown that hue and shape are "separable" dimensions, such that, for example, hue does not necessarily impact the perception of shape (Garner, 1974; Maddox, 1992). By contrast, hue and brightness are "integral" dimensions, such that across a wide range of tasks, hue is seen to automatically impact judgments of brightness and vice versa (Garner & Felfoldy, 1970; Maddox, 1992). In the present experiments, rather than examining dimensions that are perceptually integral we examine properties that can be perceived separately and must be bound in memory. We can thus ask whether we form bound memory representation out of perceptually distinct features. This is a different approach than that taken in the existing literature on holistic representations of real-world objects; for example, some evidence suggests that faces are represented holistically as integral units rather than as bound but ultimately independent features of eyes, noses, and mouths (e.g., Tanaka & Farah, 1993; although see Reinitz, Morrissey, & Demb, 1994, for evidence that holistic face encoding is may depend on attention at encoding).

Binding Objects to Contexts

Much of the literature examining binding with real-world objects has focused on binding objects to context. For example, in visual cognition it has been found that scene context can function as a retrieval cue for object details (Hollingworth, 2006); that memory for the spatial position of objects in scenes is better when the scene is presented during testing (Hollingworth, 2007; Mandler & Johnson, 1976); and that memory for object details and memory for the scene viewpoint are stored independently, rather than as a bound unit in memory (Varakin & Loschky, 2010). More broadly, long-term memory for individual items and objects is generally found to be independent of memory for the associations between items and the associations between items and contexts (e.g., Ceraso, Kourtzi, & Ray, 1998; Johnson & Raye, 2000; Marshuetz, 2005; Mather, 2007). In fact, many models of memory suppose

---

[1] For an analogous idea that encoding factors such as attention may cause the same items to be remembered in both implicit and explicit memory, see Turk-Browne, Yi, and Chun (2006).

that the hippocampus and prefrontal cortex are critically involved only in the "binding" aspects of memory important for remembering associations between various elements of an event and are not involved in memory for individual objects (Davachi & Wagner, 2002; Mitchell, Johnson, Raye, & D'Esposito, 2000; Ranganath, Cohen, Dam, & D'Esposito, 2004). These theories treat memory for objects as holistic and memory for events as requiring binding between disparate elements to form a true episodic memory.

In the current work, we find that even memory for individual objects—often used as "items" in such memory studies—is not holistic and instead that separate visual and semantic properties of objects are forgotten separately. This implies that the recognition of a real-world object is not a holistic process and that association and binding between separate visual and semantic properties are required for an object to be entirely remembered. Depending on the stimuli used, many experiments that claim to be isolating a binding mechanism by contrasting memory for objects with memory for the context in which such objects were seen may be failing to do so, as even their nonbinding condition may depend critically on binding processes within objects (see Davachi, 2006, for a discussion of within- vs. between-object binding and the role of hippocampus). Although there are likely differences between within-object binding and across-object binding (e.g., emotion seems to differentially impact these processes; Mather, 2007), the role of binding for features within real-world objects is critical to the interpretation of such memory studies.

Object Representations and Object Recognition

As both visual long-term memory and object recognition are thought to depend on the same high-level object representations (Palmeri & Tarr, 2008), memory errors such as those in the current data may be able to usefully inform models of object recognition by elucidating the underlying object representation. In particular, one of the object properties we use in the current experiments is object state, which we define as a change in the pose or configuration of an object's parts (see Brady et al., 2008). This aspect of object representation has rarely been addressed in the existing literature on visual memory, and it is likely to be an important component of object recognition: Many everyday objects contain movable units that affect the semantics or functional uses of an object while keeping visual information similar and not resulting in a change in identity of an object. Studying memory for changes in the configuration or pose of an object's parts is interesting because part-based representation is an important point of debate in the literature on object recognition and view-based versus more structured representations of objects (e.g., Palmeri & Tarr, 2008).

Although our findings do not directly address whether separate parts within an object are forgotten separately, the independence of state changes from color or exemplar changes lends some credence to structural models of object representation, where configurations of parts are explicitly represented independently of each other and could therefore be separately forgotten in memory (Hummel, 2000). Such independence fails to support view-based theories, which tend to assume holistic object representations (although see Ullman, 2007), and theories in which visual recognition is thought to proceed by increasingly complex conjunctions forming new features until an entire object is represented, which also tend to assume holistic object representations (DiCarlo & Cox, 2007; Serre, Wolfe, Bileschi, Riesenhuber, & Poggio, 2007).

Conclusion

So what is the format of real-world object representations? We find independent forgetting of information about an object's color, information that distinguishes different exemplars of the same category, and information that distinguishes changes in object state. This suggests that the underlying visual features that we rely on to distinguish these different changes are distinct and are forgotten separately. These results demonstrate that real-world objects are not represented as a single bound unit in visual memory. Furthermore, although there is often a dependence between how likely observers are to remember different properties of the same object, this appears to be due to encoding factors rather than to reflect a bound underlying memory representation.

References

Alvarez, G. A., Konkle, T., Brady, T. F., Gill, J., & Oliva, A. (2009). Comparing the fidelity of perception, short-term memory and long-term memory: Evidence for highly detailed long-term memory representations. *Journal of Vision, 9*(8), Article 584. doi:10.1167/9.8.584

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

Bays, P. M., Wu, E. Y., & Husain, M. (2011). Storage and binding of object features in visual working memory. *Neuropsychologia, 49,* 1622–1631. doi:10.1016/j.neuropsychologia.2010.12.023

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), Article 4. doi:10.1167/11.5.4

Brady, T. F., Konkle, T., Alvarez, G. A., & Oliva, A. (2008). Visual long-term memory has a massive storage capacity for object details. *Proceedings of the National Academy of Sciences, USA, 105,* 14325–14329. doi:10.1073/pnas.0803390105

Brady, T. F., Konkle, T., Gill, J., Oliva, A., & Alvarez, G. A. (2012). *Long-term memory has the same limit on fidelity as working memory.* Manuscript submitted for publication.

Bülthoff, H. H., & Edelman, S. (1992). Psychophysical support for a two-dimensional view interpolation theory of object recognition. *Proceedings of the National Academy of Sciences, USA, 89,* 60–64. doi:10.1073/pnas.89.1.60

Ceraso, J., Kourtzi, Z., & Ray, S. (1998). The integration of object properties. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 24,* 1152–1161. doi:10.1037/0278-7393.24.5.1152

Chun, M. M. (2011). Visual working memory as visual attention sustained internally over time. *Neuropsychologia,* 1407–1409. doi:10.1016/j.neuropsychologia.2011.01.029

Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences, 24,* 87–114. doi:10.1017/S0140525X01003922

Davachi, L. (2006). Item, context and relational episodic encoding in humans. *Current Opinion in Neurobiology, 16,* 693–700. doi:10.1016/j.conb.2006.10.012

Davachi, L., & Wagner, A. D. (2002). Hippocampal contributions to episodic encoding: Insights from relational and item-based learning. *Journal of Neurophysiology, 88,* 982–990.

Diana, R. A., Yonelinas, A. P., & Ranganath, C. (2007). Imaging recollection and familiarity in the medial temporal lobe: A three-component model. *Trends in Cognitive Sciences, 11,* 379–386. doi:10.1016/j.tics.2007.08.001

DiCarlo, J. J., & Cox, D. (2007). Untangling invariant object recognition. *Trends in Cognitive Sciences, 11,* 333–341. doi:10.1016/j.tics.2007.06.010

Fougnie, D., & Alvarez, G. A. (2011). Object features fail independently in working memory: Evidence for a probabilistic feature-store model. *Journal of Vision, 11*(12), Article 3. doi:10.1167/11.12.3

Fougnie, D., Asplund, C. L., & Marois, R. (2010). What are the units of storage in visual working memory? *Journal of Vision, 10*(12), Article 27. doi:10.1167/10.12.27

Gajewski, D. A., & Brockmole, J. R. (2006). Feature bindings endure without attention: Evidence from an explicit recall task. *Psychonomic Bulletin & Review, 13,* 581–587.

Garner, W. R. (1974). *The processing of information and structure*. Hillsdale, NJ: Erlbaum.

Garner, W. R., & Felfoldy, G. L. (1970). Integrality of stimulus dimensions in various types of information processing. *Cognitive Psychology, 1,* 225–241.

Hanna, A., & Remington, R. (1996). The representation of color and form in long-term memory. *Memory & Cognition, 24,* 322–330. doi:10.3758/BF03213296

Hollingworth, A. (2006). Scene and position specificity in visual memory for objects. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 32,* 58–69. doi:10.1037/0278-7393.32.1.58

Hollingworth, A. (2007). Object-position binding in visual memory for natural scenes and object arrays. *Journal of Experimental Psychology: Human Perception and Performance, 33,* 31–47. doi:10.1037/0096-1523.33.1.31

Hummel, J. E. (2000). Where view-based theories break down: The role of structure in shape perception and object recognition. In E. Dietrich & A. Markman (Eds.), *Cognitive dynamics: Conceptual change in humans and machines* (pp. 157–185). Hillsdale, NJ: Erlbaum.

Irwin, R. J., & Hautus, M. J. (1997). Likelihood-ratio decision strategy for independent observations in the same–different task: An approximation to the detection-theoretic model. *Perception & Psychophysics, 59,* 313–316.

Johnson, M. K., & Raye, C. L. (2000). Cognitive and brain mechanisms of false memories and beliefs. In D. L. Schacter & E. Scarry (Eds.), *Memory, brain, and belief* (pp. 35–86). Cambridge, MA: Harvard University Press.

Konkle, T., Brady, T. F., Alvarez, G. A., & Oliva, A. (2010). Conceptual distinctiveness supports detailed visual long-term memory for real-world objects. *Journal of Experimental Psychology: General, 139,* 558–578. doi:10.1037/a0019165

Luck, S. J., & Vogel, E. K. (1997, November 20). The capacity of visual working memory for features and conjunctions. *Nature, 390,* 279–281. doi:10.1038/36846

Luria, T., & Vogel, E. K. (2011). Shape-color and color-color conjunction stimuli are represented as bound objects in visual working memory. *Neuropsychologia, 49,* 1632–1639.

Macmillan, N. A., & Creelman, C. D. (2005). *Detection theory: A user's guide* (2nd ed.). Mahwah, NJ: Erlbaum.

Maddox, W. T. (1992). Perceptual and decisional separabilty. In F. G. Ashby (Ed.), *Multidimensional models of perception and cognition* (pp. 147–180). Hillsdale, NJ: Erlbaum.

Magnussen, S. (2000). Low-level memory processes in vision. *Trends in Neurosciences, 23,* 247–251. doi:10.1016/S0166-2236(00)01569-1

Mandler, J. M., & Johnson, N. S. (1976). Some of the thousand words a picture is worth. *Journal of Experimental Psychology: Human Learning and Memory, 2,* 529–540. doi:10.1037/0278-7393.2.5.529

Marshuetz, C. (2005). Order information in working memory: An integrative review of evidence from brain and behavior. *Psychological Bulletin, 131,* 323–339. doi:10.1037/0033-2909.131.3.323

Mather, M. (2007). Emotional arousal and memory binding: An object-based framework. *Perspectives on Psychological Science, 2,* 33–52. doi:10.1111/j.1745-6916.2007.00028.x

Maunsell, J. H. R., & Treue, S. (2006). Feature-based attention in visual cortex. *Trends in Neurosciences, 29,* 317–322. doi:10.1016/j.tins.2006.04.001

Mitchell, K. J., Johnson, M. K., Raye, C. L., & D'Esposito, M. (2000). fMRI evidence of age-related hippocampal dysfunction in feature binding in working memory. *Cognitive Brain Research, 10,* 197–206. doi:10.1016/S0926-6410(00)00029-X

Naor-Raz, G., Tarr, M. J., & Kersten, D. (2003). Is color an intrinsic property of object representation? *Perception, 32,* 667–680. doi:10.1068/p5050

O'Craven, K. M., Downing, P. E., & Kanwisher, N. (1999, October 7). fMRI evidence for objects as the units of attentional selection. *Nature, 401,* 584–587. doi:10.1038/44134

Olson, I. R., & Jiang, Y. (2002). Is visual short-term memory object based? Rejection of the "strong-object" hypothesis. *Perception & Psychophysics, 64,* 1055–1067. doi:10.3758/BF03194756

Palmeri, T. J., & Tarr, M. (2008). Visual object perception and long-term memory. In S. Luck & A. Hollingworth (Eds.), *Visual memory* (pp. 163–207). Oxford, England: Oxford University Press.

Parks, C. M., & Yonelinas, A. P. (2007). Moving beyond pure signal-detection models: Comment on Wixted (2007). *Psychological Review, 114,* 188–201. doi:10.1037/0033-295X.114.1.188

Price, C. J., & Humphreys, G. W. (1989). The effects of surface detail on object categorization and naming. *Quarterly Journal of Experimental Psychology: Human Experimental Psychology, 41*(A), 797–827. doi:10.1080/14640748908402394

Ranganath, C., Cohen, M. X., Dam, C., & D'Esposito, M. (2004). Inferior temporal, prefrontal, and hippocampal contributions to visual working memory maintenance and associative memory retrieval. *Journal of Neuroscience, 24,* 3917–3925. doi:10.1523/JNEUROSCI.5053-03.2004

Ranganath, C., Yonelinas, A. P., Cohen, M. X., Dy, C. J., Tom, S. M., & D'Esposito, M. (2004). Dissociable correlates of recollection and familiarity within the medial temporal lobes. *Neuropsychologia, 42,* 2–13. doi:10.1016/j.neuropsychologia.2003.07.006

Reinitz, M. T., Morrissey, J., & Demb, J. (1994). Role of attention in face encoding. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 20,* 161–168. doi:10.1037/0278-7393.20.1.161

Rensink, R. A., O'Regan, J. K., & Clark, J. J. (1997). To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science, 8,* 368–373. doi:10.1111/j.1467-9280.1997.tb00427.x

Riesenhuber, M., & Poggio, T. (1999). Hierarchical models of object recognition in cortex. *Nature Neuroscience, 2,* 1019–1025. doi:10.1038/14819

Riesenhuber, M., & Poggio, T. (2000). Models of object recognition. *Nature Neuroscience, 3,* 1199–1204. doi:10.1038/81479

Scholl, B. J. (2001). Objects and attention: The state of the art. *Cognition, 80,* 1–46. doi:10.1016/S0010-0277(00)00152-9

Serre, T., Wolf, L., Bileschi, S., Riesenhuber, M., & Poggio, T. (2007). Robust object recognition with cortex-like mechanisms. *IEEE Transactions on Pattern Analysis and Machine Intelligence, 29,* 411–426. doi:10.1109/TPAMI.2007.56

Stefurak, D. L., & Boynton, R. M. (1986). Independence of memory for categorically different colors and shapes. *Perception & Psychophysics, 39,* 164–174.

Tanaka, J., & Bunosky, L. (1993, November). *Is the object recognition system really colorblind?* Poster session presented at the meeting of the Psychonomic Society, Washington, DC.

Tanaka, J. W., & Farah, M. J. (1993). Parts and wholes in face recognition. *Quarterly Journal of Experimental Psychology, 46,* 225–245.

Tarr, M. J., & Bülthoff, H. H. (1995). Is human object recognition better described by geon structural descriptions or by multiple views? Comment on Biederman and Gerhardstein (1993). *Journal of Experimental Psychology: Human Perception and Performance, 21,* 1494–1505. doi:10.1037/0096-1523.21.6.1494

Treisman, A. (1998). Feature binding, attention and object perception. *Philosophical Transactions of the Royal Society of London, Series B: Biological Sciences, 353,* 1295–1306. doi:10.1098/rstb.1998.0284

Tulving, E., & Thomson, D. M. (1973). Encoding specificity and retrieval processes in episodic memory. *Psychological Review, 80,* 352–373. doi:10.1037/h0020071

Turk-Browne, N. B., Yi, D. J., & Chun, M. M. (2006). Linking implicit and explicit memory: Common encoding factors and shared representations. *Neuron, 49,* 917–927.

Ullman, S. (2007). Object recognition and segmentation by a fragment-based hierarchy. *Trends in Cognitive Sciences, 11,* 58–64.

Varakin, D. A., & Loschky, L. (2010). Object appearance and picture-specific viewpoint are not integrated in long-term memory. *Quarterly Journal of Experimental Psychology, 63,* 1181–1200. doi:10.1080/17470210903254639

Vogel, E. K., Woodman, G. F., & Luck, S. J. (2001). Storage of features, conjunctions, and objects in visual working memory. *Journal of Experimental Psychology: Human Perception and Performance, 27,* 92–114. doi:10.1037/0096-1523.27.1.92

Wheeler, M. E., & Treisman, A. M. (2002). Binding in short-term visual memory. *Journal of Experimental Psychology: General, 131,* 48–64. doi:10.1037/0096-3445.131.1.48

Wiseman, S., & Neisser, U. (1974). Perceptual organization as a determinant of visual recognition memory. *American Journal of Psychology, 87,* 675–681. doi:10.2307/1421974

Wixted, J. T. (2007). Dual-process theory and signal-detection theory of recognition memory. *Psychological Review, 114,* 152–176. doi:10.1037/0033-295X.114.1.152

Wixted, J. T., & Stretch, V. (2000). The case against a criterion-shift account of false memory. *Psychological Review, 107,* 368–376.

Xu, Y. (2002). Limitations in object-based feature encoding in visual short-term memory. *Journal of Experimental Psychology: Human Perception and Performance, 28,* 458–468. doi:10.1037/0096-1523.28.2.458

Yonelinas, A. P. (2002). The nature of recollection and familiarity: A review of 30 years of research. *Journal of Memory and Language, 46,* 441–517.

Appendix A

Binding Model Derivation

Here we derive specific predictions for the dependence between state and exemplar accuracy given a fully bound model, dependence(D) = 1, correcting for the fact that guessing is independent between features by definition.

If memory representations are bound, the probability of getting the state correct given you get the exemplar correct will be depend on whether you remember the item or not. In general:

$$p_{D=1}(state \mid exemplar) = \frac{p(state \cap exemplar)}{p(exemplar)}$$

$$p_{D=1}(state \mid exemplar) = \frac{R(pc) + \frac{1}{4}G(pc)}{R(pc) + \frac{1}{2}G(pc)}$$

On the other hand, observers should get the state comparison correct when they fail to get exemplar comparison correct only if they guess it correctly, because the bounded model posits the two memories are never recalled independently:

$$p_{D=1}(state \mid \sim exemplar) = \frac{p(state \cap \sim exemplar)}{p(\sim exemplar)}$$

$$p_{D=1}(state \mid \sim exemplar) = \frac{\frac{1}{4}G(pc)}{\frac{1}{2}G(pc)}$$

Correspondingly, $p^{+}(state \mid exemplar)$, the amount more likely observers are to get state correct if they get exemplar correct than if they get exemplar incorrect, is the subtraction of these two terms:

$$p_{D=1}^{+}(state \mid exemplar) = \frac{R(pc) + \frac{1}{4}G(pc)}{R(pc) + \frac{1}{2}G(pc)} - \frac{\frac{1}{4}G(pc)}{\frac{1}{2}G(pc)}$$

Which reduces to simply

$$p_{D=1}^{+}(state \mid exemplar) = \frac{R(pc)}{R(pc) + 1}$$

The following is Matlab code to simulate the bound model using Monte Carlo methods:

(Appendices continue)

```
nSubjects = 15;
nItems = 120;

percentCorrList = 0.5:0.01:1.0;
for percI = 1:length(percentCorrList)
  percentCorr = percentCorrList(percI);

  % Do 500 simulations at each percent correct
  for m=1:500

    % What percentage of items should we remember?
    R_mean =  2*percentCorr - 1;

    % Generate simulated data with that number of items remembered
    R = rand(nSubjects, nItems)<=R_mean;

    % Generate independent guesses for state and exemplar
    G_S = rand(nSubjects, nItems)<=0.50;
    G_E = rand(nSubjects, nItems)<=0.50;

    % We get a comparison correct if we remember it or guess correctly
    stateCorrect = R | G_S;
    exempCorrect = R | G_E;

    % Save condition probability (increased likelihood of getting
    % state given you also get exemplar):
    condProb(m, percI) = mean(stateCorrect(exempCorrect==1))...
          - mean(stateCorrect(exempCorrect==0));
  end
end

% Plot function
plot(percentCorrList, mean(condProb));
```

Appendix B

How Independent Is Forgetting?

To quantify the degree of independence between the underlying memory representations for state and exemplar information, we asked what percentage of the forgetting between our short- and long-delay conditions appears to be independent forgetting rather than correlated, bound forgetting. In particular, we model memory (a) representations that are truly independent and are thus forgotten independently but are correlated at initial encoding and (b) representations that are bound and are thus always forgotten together. We then estimate what proportion of bound versus independent representations is needed to best explain the data from both the short- and long-delay conditions.

Critically, we find the combination of feature boundedness and encoding correlations that simultaneously fit both the short- and long-delay conditions requires nearly 100% independent forgetting of the features (almost no truly bound representations). Pseudocode (in the style of Matlab code) for this simulation is presented below:

```
% PARAMETERS:
% -----------
percentCorrectAtShortDelay = 0.72;
lossInPercentCorrectWithDelay = 0.09;

initialCorrelation = 0.65;
amountTrulyBound = 0.03;
```

(Appendices continue)

```
% SHORT DELAY:
% ------------

% Sample initially correlated, yet ultimately independent state
% and exemplar memories (of size [nSubs, nItems])
[Smem, Emem] = CreateCorrelatedUniforms(initialCorrelation,
[nSubs,nItems]);

S_mean =  2*percentCorrectAtShortDelay-1;
S = Smem <= S_mean;

E_mean =  2*percentCorrectAtShortDelay-1;
E = Emem <= E_mean;

% Simulate bound memories
Rmem = rand(nSubs,nItems);
R_mean =  2*percentCorrectAtShortDelay-1;
R = Rmem <= R_mean;

% Compute independent guesses for state and exemplar
G_S = rand(nSubs,nItems)<=0.50;
G_E = rand(nSubs,nItems)<=0.50;

% Mix trials from bound and independent memories in appropriate ratio:
useBoundForTrial = rand(nSubs,nItems) <= amountTrulyBound;

% And calculate overall percent correct for state & exemplar:
stateCorrect = (~useBoundForTrial & S) | (useBoundForTrial & R) | G_S;
exempCorrect = (~useBoundForTrial & E) | (useBoundForTrial & R) | G_E;

% Now calculate conditional probability for short delay:
condProb_ShortDelay = mean(stateCorrect(exempCorrect==1)) ...
  - mean(stateCorrect(exempCorrect==0));


% LONG DELAY:
% ------------

% Now forget some memories -- forget independent memories
independently:
numItemsToForget = round(nItems*(lossInPercentCorrectWithDelay));
for j=1:size(S,1)
  available = Shuffle(find(S(j,:)==1));
  S(j,available(1:numItemsToForget)) = 0;
end

for j=1:size(E,1)
  available = Shuffle(find(E(j,:)==1));
  E(j,available(1:numItemsToForget)) = 0;
end

% ... and forget bound memories in a bound fashion:
for j=1:size(R,1)
  available = Shuffle(find(R(j,:)==1));
  R(j,available(1:numItemsToForget)) = 0;
end

% Now calculate percent correct/condProb again:
G_S = rand(subs,items)<=0.50;
G_E = rand(subs,items)<=0.50;

stateCorrect = (~useBoundForTrial & S) | (useBoundForTrial & R) | G_S;
exempCorrect = (~useBoundForTrial & E) | (useBoundForTrial & R) | G_E;

% Now calculate conditional probability for long delay:
condProb_LongDelay = mean(stateCorrect(exempCorrect==1)) ...
  - mean(stateCorrect(exempCorrect==0));
```

Using this model we can compute predictions for each combination of boundedness and encoding correlation, given the percent correct we observe for short and long delay. We can then compare the condition probability predicted by those models to that we actually observe and compute an error term (root mean square error). These raw errors are plotted in Figure B1; the minimum error values are plotted in Figure 5 in the main text.

(Appendices continue)

Figure B1. Best fit models that vary in both the memory boundedness and the correlation between state and exemplar encoding. (a, b) Both the short-delay and long-delay conditions can be fit by assuming a correlation at encoding, a bounded memory representation, or any mixture of the two. (c) However, the combined data can be fit only by assuming that nearly all of the forgetting is independent (i.e., that the dependence results almost entirely from correlations at encoding). This is because observers' dependence decreases much more rapidly than would be expected by a model with bound memory representations. The best fit parameters for explaining both the short-delay and long-delay data suggest 97% independence between state and exemplar properties with a correlation of $r = .63$ in the initial encoding of state and exemplar properties. RMSE = root-mean-square error.

Appendix C

Comparison of High-Threshold and Signal Detection Models

Our model of how bound a memory representation appears is based on a high-threshold memory model. In particular, we assume that observers successfully remember some proportion of the items and have no information about others. Such high-threshold models provide reasonable fits to recollection data (Parks & Yonelinas, 2007; Yonelinas, 2002) and thus should be sufficient to distinguish between our bound and independent hypotheses, even if ultimately signal detection may be a better model of recollection processes (e.g., Wixted, 2007). In addition, forced-choice tasks in general are less sensitive to the distinction between signal detection and threshold models because, when there is little response bias, there is little dependence on the particular shape of the response operating characteristic curve that distinguishes these models (Macmillan & Creelman, 2005). Although these points mitigate concerns over using a high-threshold model, it is also possible to examine our data using signal detection.

In particular, we can model a bound hypothesis as reflecting a single underlying memory signal. Specifically, the correct item's memory strength would be reflected by a normal distribution centered at d′, with the three distractors each centered around 0. This framing of our task is in line with a signal detection model of the Deese–Roediger–McDermott task (e.g., Macmillan & Creelman, 2005; Wixted & Stretch, 2000). The independence hypothesis, by contrast, would be modeled as reflecting two underlying memory signals: a state familiarity signal and an exemplar familiarity signal. Thus, the underlying memory signal must be considered in a two-dimensional space, where the correct item's memory signal is reflected by a normal distribution centered at (state d′, exemplar d′) and the distractors are centered at (state d′, 0), (exemplar d′, 0), and (0,0).

To model performance in our task, we must convert these underlying memory strengths into a model of our particular four-alternative forced-choice comparison. In the case of a straightforward two-alternative forced-choice task with unrelated targets and foils, $d' = \sqrt{2} \times z(\text{pc})$, where z is the inverse cumulative normal distribution function (Macmillan & Creelman, 2005). This is because the distance between two orthogonal normal distributions, each centered at d′, is $\sqrt{2} \times d'$. To generalize to our higher dimensional stimuli and four-alternative task, we can simulate the process that leads to this formula by using Monte Carlo methods. In particular, in the two-alternative forced-choice case, we can sample a large number of memory strengths from a normal distribution centered at d′ and from one centered at 0 and, for each pair, ask how likely the greater memory strength is to be from the correct item, rather than the foil. To generalize to the four-alternative forced-choice case, we can simply sample from the memory strengths of each of the four items and once again ask how likely the item with the highest memory strength is to be correct on state and/or exemplar to determine a percent correct.

(Appendices continue)

Figure C1. (a). Results if the independence model relies upon summed familiarity choice rule. (b) Results if independence model relies upon likelihood ratio choice rule. In each plot, the different lines correspond to different underlying pc/d' in the simulated data. The dashed black line corresponds to equality, $x = y$; pc = percent correct.

For the independent model, a decision must be made about how observers pool information from the two memory signals to choose a single answer. There are at least two possibilities: (a) they choose the item with the largest summed familiarity signal; (b) they choose the item whose combined memory signal is most likely to have come from an "old" item (e.g., according to the likelihood ratio; Irwin & Hautus, 1997). In the current model, this means the item most likely to have been generated by a normal distribution centered at (state d', exemplar d').

Finally, we can take these modeled four-alternative forced choices and ask, if we fit our high-threshold model to these data, how bounded does the result look? In particular, we can mix samples from the bound model and the independent model together in a certain ratio, simulating partially bound memory representations, and ask whether our high-threshold model can accu-

rately recover the percent bounded that is simulated according to the signal detection model. We find that it can (see Figure C1). Although the recovered boundedness scores do systematically deviate from the modeled boundedness, they do so in a well-behaved, linear fashion that differs little based on the underlying percent correct/d' of the model. Thus, even if signal detection is a better model of the underlying memory traces and decision process, our conclusions remain unaffected: The dependence of state and exemplar decreases systematically over time and does so at a rate much greater than we would expect by the decrease in percent correct or, correspondingly, d'.