The Case of the Missing Visual Details: Occlusion and Long-Term Visual Memory

Carrick C. Williams and Kyle A. Burkle

California State University San Marcos

Author Note

Carrick C. Williams, Psychology Department, California State University San Marcos; Kyle A. Burkle, Psychology Department, California State University San Marcos.

Portions of the data were reported at the 56th annual meeting of the Psychonomic Society, Chicago, IL, November 2015 and the 57th annual meeting of the Psychonomic Society, Boston, MA, November 2016. We thank Alisha Kent, Conner Maloney, Angie Rodriguez, Danielle Struwe, and Anahi Guerrero for their assistance in data collection. Correspondence concerning this article should be addressed to Carrick Williams, Department of Psychology, 333 S Twin Oaks Valley Rd, San Marcos, CA 92096. Email: cawilliams@csusm.edu.

10.1037/xlm0000393

Abstract

To investigate the critical information in long-term visual memory representations of objects, we used occlusion to emphasize one type of information or another. By occluding one solid side of the object (e.g., top 50%) or by occluding 50% of the object with stripes (like a picket fence), we emphasized visible information about the object, processing the visible details in the former and the object's overall form in the latter. On a token discrimination test, surprisingly, memory for solid or stripe occluded objects at either encoding (Experiment 1) or test (Experiment 2) was the same. In contrast, when occluded objects matched at encoding and test (Experiment 3) or when the occlusion shifted, revealing the entire object piecemeal (Experiment 4), memory was better for solid compared to stripe occluded objects, indicating that objects are represented differently in long-term visual memory. Critically, we also found that when the task emphasized remembering exactly what was shown, memory performance in the more detailed solid occlusion condition exceeded that in the stripe condition (Experiment 5). However, when the task emphasized the whole object form, memory was better in the stripe condition (Experiment 6) than in the solid condition. We argue that long-term visual memory can represent objects flexibly, and task demands can interact with visual information, allowing the viewer to cope with changing real-world visual environments.

Keywords: Long-Term Visual Memory; Occlusion; Object Memory; Picture Representation

When viewing a complex scene, such as a street scene, many objects will be partially visible from a single point of view. Researchers who have explored the impact of occlusion on object perception have primarily focused on the ability of observers to perceive complete objects behind an occluding element (amodal completion; e.g., Kellman & Shipley, 1992; Murray, Sekuler, & Bennett, 2001; Palmer, Kellman, & Shipley, 2006; Rensink & Enns, 1998; Sekuler & Palmer, 1992). In other words, consistent with Gestalt principles, people appear to fill in the missing components when looking at an object. Although there may be questions regarding the ability to use amodal completion in guiding attention (Wolfe, Reijnen, Horowitz, Pedersini, Pinto, & Hulleman, 2011), it is clear that individuals can identify an object even if it extends behind another object. The ability to identify occluded objects and potentially "fill in" the missing components indicates that people can easily deal with this real-world viewing problem in real-time. We explored if these abilities had consequences for the representations of objects in long-term visual memory.

The question of how occluded objects are represented may appear to be trivial given that they are so easily dealt with in everyday life, but how one represents an occluded object could be informative regarding the limitations and flexibility of long-term visual memory representations. Objects are stored in long-term visual memory with a great amount of detail (e.g., Brady, Konkle, Alavarez, & Oliva, 2008; Standing, 1973). With regard to memories of occluded objects, it is possible that the number of visual features that are disrupted by the occlusion could impact the ability to represent an object effectively. In addition, even if one "fills in the gaps" in a picture during an online viewing task (via a process like amodal completion), it is possible that long-term visual memory may not obligatorily represent the object the same way. Long-term visual memory appears to rely on conceptual information about the object's basic and subordinate category structure (Konkle, Brady, Alvarez, & Oliva, 2010), and this conceptual information can be modified by task demands (Antonelli & Williams, 2016). Long-term visual memories could be more tied to encoding demands of the task and thus, different visual memory representations could be created that either represent an incomplete object or fill in the missing components in some manner.

Related to the importance of conceptual information to long-term visual memories, occlusion could obscure a feature of the object that identifies either the semantic category or the subcategory of the object. For example, silverware handles can be very similar or identical. If the other end of a utensil is occluded by a napkin, it may be impossible to know whether the specific object is a spoon or a fork. Although one could change his or her point of view to see the critical feature (or simply move the occluding element), it is not always possible. If the defining characteristic of the category of an object is occluded, the ability to store and remember that object would likely be impaired because one could not use the critical conceptual information about the object's category to support visual memory (as in Konkle et al. 2010).

When confronted with visual information that is incomplete, how does the long-term visual memory system represent the object? Instead of point-by-point veridical representations of the real-world, long-term visual memories have been described as abstract visual representations (Hollingworth & Henderson, 2002). Although they are abstract in the sense of not being exactly what was seen, these long-term visual memories are detailed enough that they can be used to perform a number of subtle discriminations, such as mirror reversals (Castelhano & Henderson, 2005). Long-term visual memory representations have also been shown to possess a massive capacity for visual details (e.g., Brady et al., 2008; Standing, 1973). Following this description, we are not arguing for a memory representation that is equivalent to the exact pixels of the object, but rather a more abstract representation that is stored in long-term visual memory and can be retrieved. Although these representations are abstract, there is a question of how abstract they are. Because these memories are not likely veridical representations, we wanted to examine the precision of the representations in long-term visual memory by exploring how the long-term visual memory system accounts for the common situation of incomplete/missing visual details.

The first way that long-term visual memory could address a lack of visual information due to occlusion would be to simply represent only the visible portions of the object, ignoring the occluding element. This type of representation would be a more veridical way of representing the object as it was viewed. If one previously saw an apple that was blocked from view by a bowl, this type of representation of the apple would only contain the visible portions. In other words, a partial, but identifiable and categorizable, apple is stored in long-term visual memory. The memory representation would contain substantial holes, but from the point of view of what was actually seen, it would be a relatively accurate representation. Again, we are not arguing for a retinotopic or similar level representation in long-term visual memory, but this manner of dealing with occluded objects would represent what was and was not seen from a viewpoint.

A second way that the memory system could address occlusion is that the missing information is "filled in" prior to the memory being stored (i.e., amodally completed). In other words, the representation of the object contains the viewer's inference of what the remainder of the object looks like. Unlike a representation retaining details about what was occluded or not, the filled in representation could be a generalized abstract form that glosses over the presented details (similar to a geon structural description, Biederman, 1987). Alternatively, it could contain areas of high detail (i.e., the unoccluded areas) and areas of low detail (i.e., the filled in areas) in a sort of patchwork memory. Regardless of the form, the representation necessarily would be unlike the originally viewed object because the inferences made to complete the object (either explicitly or implicitly) would be based on top-down information and would be included in the memory representation. Potentially, this inference process would create an even more abstract representation of an occluded object (i.e., a representation that is less tied to bottom-up information of the image, Hollingworth & Henderson, 2002) than a representation of a fully visible object because the inferences made could be more generic/prototypical than the actual image. Similar to the concept of boundary extension (Intraub & Richardson, 1989), in which scenes are remembered from a wider angle than they were viewed (i.e., containing more information than was originally presented), the abstract representation of the objects would be inaccurate from the point of view of what was exactly seen, but useful in identifying what was seen generally.

To test these two possible ways that partial objects could be represented, we used two forms of occlusion, which we called ***solid*** and ***stripe****.* The solid occlusion blocked an entire half of the object (e.g., the left half) leaving the other half fully visible. The solid occlusion is similar to the apple and bowl example given above. By leaving an entire side visible, observers could see all of the visual details of that side of the object, but they had no information about the occluded side. With images that are simple or symmetrical, the fact that one side of the object was missing may not be difficult to overcome. However, with more complex objects (e.g., side views of airplanes or tractors), the occlusion of an entire half of the object could be detrimental. Regardless of symmetry, the object that is seen with one side occluded allows for the more extensive processing of the visible portion and may result in a representation of the visual details of that side only.

On the other hand, the stripe occlusion likely encourages processing of the whole form of the object with less emphasis on the details of any one part. The objects were presented in alternating visible and masked stripes. The same total area (50%) of the object was occluded, but the object appeared as if it were behind a picket fence (when the stripes were vertical) or a set of window blinds (when the stripes were horizontal). By occluding the object with stripes, the extent of the whole object shape could be processed, but the visual information would be discontinuous. The use of stripes is similar to Henderson's (1997) study of transaccadic memory for objects. Henderson presented the object's visual information with alternating stripes prior to the eye movement (when the object was in the periphery) and after the eye movement (when the object was fixated). He found that participants were as quick at responding if the stripes shifted during an eye movement as when the stripes did not move. In other words, participants were insensitive to a change that completely altered the visual information held in their memory. Henderson interpreted these results as evidence that transsaccadic memory relied on the general structural description of the object rather than the precise visual information that was extracted. Similarly, Henderson and Hollingworth (2003) found that participants were insensitive to changes to stripe occluded scenes across eye movements (in this case, every pixel of the image changed when the eyes crossed invisible boundaries and participants detected fewer than 6% of the changes in Experiment 1), indicating that participants were creating an abstract visual representation. Although these previous results were testing transsaccadic memory rather than long-term visual memory representations, the idea of the using stripe occlusion as a test for whole object form representations motivated our use of this type of occlusion.

Although how long-term visual memory deals with occlusion is an interesting question, it can also speak to what can and cannot be stored in long-term visual memory more generally. A memory system that “fills in the gaps” would be one that emphasizes the general form of the object in long-term visual memory. Because the information that would be “filled in” would have to be inferred by the viewer, it would not contain the specific visual details that had been occluded in the original image. By filling in the missing components and storing the result, a whole object would be represented in visual memory. When asked to retrieve the information, even when confronted with a fully visible object at test, the overall form would be the most important component to match to the representation (as in Henderson, 1997). In contrast, if long-term visual memory can represent incomplete objects, it would indicate that the visual details that are tied to what was seen are the critical pieces of information that need to be stored and the overall form is less important. In other words, the memory system would be relying on the visual features that are present to represent the object (see Hollingworth & Franconeri, 2009, for a similar argument in object tracking across occlusion).

In the current study, we briefly presented either fully visible objects (our upper limit baseline) or objects that had 50% of the visible area of the object occluded with either a stripe or solid mask. Better memory for the stripe condition compared to the solid condition would indicate that abstract, whole forms would be closer the normal way that objects are represented in long-term visual memory. In contrast, better memory for the solid condition compared to the stripe condition would indicate that long-term visual memories rely more on the details that were originally visible, rather than the whole form of the object.

General Method

Participants

For each of the six experiments, 72 naïve participants were recruited from the Human Participant Pool at California State University San Marcos (total N=432 participants). Participants received partial credit to fulfill a class requirement as part of their participation.

Design

In order to investigate the effects of occlusion on memory, stimuli could be occluded at encoding, the memory test, both, or neither (See Table 1 for list of manipulations). In each experiment, participants were presented with three blocks of trials with each block having an encoding phase followed by memory test phase. For all experiments, each encoding-memory test block consisted of trials for only one occlusion condition. In other words, a participant would encode all of the objects with one type of occlusion (e.g., solid occlusion) and then be tested on those objects before being presented with a second occlusion condition (e.g. stripe occlusion) and then a third condition (e.g., fully visible). Occlusion condition (Solid, Stripe, and Fully Visible) was manipulated within participants; the occlusion condition order was counterbalanced across participants and was included as a between participant factor in the analysis.

Materials

Sixty-four object categories with three different colors each (e.g., silver, black, and green binoculars) were taken from Williams (2010a; 2010b) and Antonelli and Williams (2016). In each block, one exemplar from each of the 64 categories was presented; in the other blocks, the other two colors of that exemplar were presented (counterbalanced across participants). Each exemplar was a unique picture and was presented at encoding for one second. Each exemplar was 90 pixels along its longer dimension and was centered vertically and horizontally on a 100 x 100 pixel background with a neutral gray color (RGB = 120). The screen resolution was set to 800 x 600 pixels and the 100 x 100 pixel area of the object subtended 5.1° of visual angle on a side when viewed at a distance of 57 cm (viewing distance was not controlled in these experiments), approximately the size of the radius of the parafovea (Rayner, 2009).

For objects that were occluded, 50% of the 100 x 100 pixel area of the object was blocked from view by a multi-colored mask created with Adobe Photoshop®'s add noise filter set at the maximum noise level (400%). Two different occlusion patterns, Solid and Stripe, were used along with a fully visible condition for all experiments (see Figure 1). In the *Solid* occlusion condition, 50% of the visible area of the object was occluded using the multi-colored mask (right half, left half, top half, or bottom half equally distributed across the 64 objects in that condition). The mask was opaque, completely blocking the view of half of the object; the other half was visible. In the *Stripe* occlusion condition, objects were again presented with 50% of the visible area of the objects occluded, but the occlusion was formed by multi-colored bars that completely occluded the object information alternating with fully visible stripes of the objects (either horizontal or vertical stripes). These bars create a kind of "picket fence" view of the object (Henderson, 1997). A total of five multi-colored 10 x 100 pixel bars (approximately .5° x 5.1°) were overlaid on the object image (the solid condition used the same five bars, but without any space between them; all objects used the same masks). We created four versions of the stripe condition for each image with two horizontal stripe presentations (with the first horizontal stripe at the top being either masked or visible) and two vertical stripe presentations (with the first vertical stripe on the left being either masked or visible). As with the solid condition, the four orientations of the stripes were equally distributed across the 64 objects in that condition. In the occlusion conditions, the portion occluded was counterbalanced across participants so that every portion (top, right, left, and bottom) was occluded equally often for all objects.

In the memory test, participants were presented with two objects, one to the left and the other to the right, in a two-alternative forced choice token discrimination task (similar to Castelhano & Henderson, 2005; Williams, Henderson, & Zacks, 2005). Depending on the experiment (again see Table 1), objects could be fully visible or occluded in the memory test. One of the objects was the studied object and the other was a color-category matched foil (e.g., two different black binoculars; see Figure 1). This type of test limits the ability of participants to use categorical information to identify which object had been presented, requiring a response based on visual memory. The foil and presented objects were counterbalanced across participants.

Procedure

The Institutional Review Board of California State University San Marcos approved the procedures used here prior to beginning the study. Participants completed an informed consent form and were assigned to one of the counterbalanced condition sequences. Participants were instructed that they would see a sequence of objects presented one at a time in the center of the display followed by a memory test (forming the first block of trials) and that the series of events (sequence of objects followed by a memory test) would repeat two more times for a total of three blocks. In the instructions, they were informed that objects may or may not be completely visible and were given four examples of the locations of occlusion (bottom, top, left, and right occlusion) that may occur in that block (for the fully visible condition, four fully visible objects were presented). After indicating that they understood the encoding instructions, participants pressed a button to initiate the encoding sequence. The 64 objects for that block were presented in a random order for one second each followed by an interstimulus interval of one second with a “+” displayed in the center of the display. Following the interstimulus interval, another randomly selected object was presented. This pattern repeated until all 64 stimuli were presented (see Figure 1).

Following the encoding sequence, participants were given instructions for the memory test. They were informed that they would be presented with two objects, left and right; one of the objects was presented in the encoding sequence and the other was a matched foil. They were told to press the button on the keyboard (“z” for left and “m” for right) corresponding to the object that they had seen. Memory test stimuli were presented until a response was given; once a response was entered, the next memory test trial would appear after a blank 500 millisecond intertrial interval. Memory trials were presented in a random order until all 64 studied items had been tested for that block (across the three blocks, 192 memory test trials occurred). Participants were told to be as accurate as possible in the memory test, but to guess if they were uncertain.

Following the test phase of each series, the next encoding series began with instructions indicating the occlusion condition of the upcoming sequence. The next block contained the same 64 categories of objects, but the objects were different tokens and were in a different color than the other blocks. For example, if a participant studied an image of black binoculars in the first block, s/he could see a different image of silver binoculars in the second block, and another image of green binoculars in the third block. In total, participants were presented with 192 unique objects from 64 object categories during the encoding phases (from Williams 2010a).

The experiment was run using EPrime 1.2 software (Schneider, Eschman, & Zuccolotto, 2002) and lasted approximately 25 minutes.

Experiment 1

In Experiment 1, we used the general paradigm with participants studying occluded objects (and a fully visible baseline) followed by a fully visible memory test.

Results and Discussion

The data for each experiment were analyzed using a mixed factorial ANOVA with order of conditions being a between-participant factor and the conditions themselves being a within-participant factor. Alpha level was set at p=.05.

Not surprisingly, memory performance was affected by encoding condition, F (2, 138) = 46.69, $p < .001$, $\eta_p^2 = .404$, with memory performance being best in the fully visible encoding condition compared to either of the occlusion conditions, both ts (71) > 7.00, ps < .001 (See Table 2). More interestingly, memory performance in the solid and stripe conditions were statistically equal, t (71) = .515, p >.20. The surprisingly small difference between the two different occlusion conditions (less than 1%) indicates that the visual memory system is either insensitive to or ignores the difference in the presentation pattern of objects.

With respect to the condition order, although there was no main effect, F (2, 69) = 1.88, p = .16, $\eta_p^2 = .052$, there was a significant interaction between the occlusion condition and condition order, F (4, 138) = 7.38, $p < .001$, $\eta_p^2 = .176$. The interaction resulted from participants' memory declining from the first block to the second block to the third block for the fully visible condition, but the patterns were not as orderly for the other two conditions. Memory performance in both occlusion conditions shows a decline from the first to the third block, but the solid condition was slightly better in the second block than in the first block and the stripe condition was worst in the second block compared to all other blocks. The general pattern of memory performance dropping across blocks is consistent with proactive interference building up across the blocks. However, there may be some carryover from having seen a different pattern of occlusion in a previous block on encoding of objects with a new pattern of occlusion. To ensure that the pattern of the data that was described above for the conditions was not overly affected by the order of conditions, we also conducted an analysis of the first condition presented to the participants (i.e., Block 1) as a between participant factor, eliminating any possible carryover effects. The same pattern of memory results held in this reduced analysis with the fully visible objects (M = .87, SD = .098) being remembered better than either solid (M = .75, SD = .094) or stripe (M = .76, SD = .075) occluded objects, F (2, 71) = 13.45, p < .001; the latter two conditions, again, did not differ, t (46) = .425, p >.20.

One final analysis was performed to examine if the location of the occluding element (top, right, bottom, or left) affected the findings of similar memory performance for the stripe and solid occlusion conditions. It is possible that a solid occluder blocking the top of the object would be more disruptive than a stripe occluder that has the top stripe occluded. Because the location of the occluding element was counterbalanced, for every object, the occluder occurred in all four locations equally often across participants. The analysis indicated that, although there was a main effect of the location of the occluder, F (3, 213) = 3.36, p = .02, η_p^2 = .045, there was no interaction with occlusion condition, F < 1. We explored the main effect with pairwise comparisons and found that the only significant difference was that objects occluded on the right side were remembered better than objects occluded at the top (Bonferroni corrected p = .048). Although this last comparison is *post hoc* in nature, the worst performance for the top occluded objects indicates that more visual information may be contained in the top of the objects helping with recognition.

The purpose of this experiment was to examine if disrupting the visible portions of an object by using occlusion would degrade the encoded representation in a predictable way. The solid occlusion condition allows participants to acquire good visual information about half of the object. If having good, continuous visual information was the principle piece of information in visual memory representations, we expected that memory performance in the solid condition would be greater than memory in the stripe condition and may have approached the performance of the fully visible condition. On the other hand, if a global structural description of the object was more critical to visual memory representations, then memory in the stripe condition should have been greater than memory in the solid condition and been closer in performance to the fully visible condition. In the end, the fact that memory in the solid and stripe conditions was less than 1% different indicates that either type of information is equally useful when performing this type of memory test.

Why would two conditions that emphasize different aspects of the object be remembered equally well? The first possibility is that, instead of how the object is occluded, the important factor in visual memory is how much of the object is occluded. As a purely bottom-up explanation, this interpretation has an intuitive appeal because for both occlusion conditions, 50% of the object's visible area was occluded, and thus, they should show equivalent memory. Although this argument is intuitive, Biederman (1987) demonstrated that not all occlusion is the same in object recognition. In his study, occluded objects with recoverable geons were better recognized than objects with nonrecoverable geons. Even though the same amount of the object was occluded, participants were faster at identifying the objects in the recoverable condition. Although Biederman's claims were made about object recognition, it would seem reasonable to assume that some of the same processes would be involved in forming long-term visual memories.

A second possibility is that, even though the different types of occlusion were intended to emphasize different aspects of visual information, the participant could have extracted the same critical piece of visual information that allows him or her to remember the object. For example, if the ability to construct a structural description of the occluded object is the critical piece of visual information necessary to remember the object later, it is possible that in both the solid and stripe conditions, participants were able to extract a serviceable structural description of the object. Compared to the fully visible condition, the representation would be incomplete, but in both stripe and solid conditions the same general information is represented. In other words, even though the point of the different occlusion conditions was to focus participants on different visual information, participants were able to extract the same critical information in both conditions and use that common critical information to create the same object representation in memory.

The idea that two different patterns of occlusion resulted in ultimately the same information being encoded in long-term visual memory is similar to a critical component of Biederman's (1987) recognition by components theory, viewpoint invariance (for a contrasting argument see Tarr & Bülthoff, 1995). According to Biederman, objects should be recognized regardless of orientation as long as the geons can be extracted. The results of Experiment 1 could be interpreted in a similar manner to a viewpoint invariant system because, if participants ultimately store the same information when viewing both the stripe and solid conditions (ignoring exactly what is shown or how the object is occluded), then it would not be a surprise that later memory performance would be equal. Although an interesting extension of the recognition by components theory of Biederman, other research on long-term visual memory argues against a viewpoint invariant representation in long-term visual memory. For example, Castelhano and Henderson (2005) demonstrated that participants were able to identify the image of an object shown when tested against the mirror reversal of the image. In addition, Brady et al. (2008) pointed out that participants were able to accurately remember the "state" of the object in a two alternative choice task (e.g., a telephone with its receiver is in its cradle versus the same telephone slightly rotated with the receiver lying beside the telephone), indicating the ability to remember visual details that are specific to the particular image viewed.

A third possible explanation for equivalent memory in the two occlusion conditions is that, even though the visual information represented in the stripe and solid occlusion conditions was not the same, the information extracted in either case was sufficient to perform the memory test—a two alternative forced-choice token discrimination task. Participants may have represented objects in the stripe and solid conditions differently, as intended, but the visual memory test we used allowed the participants to use either the half object representation (solid condition) or the overall object form representation (stripe condition) equally well. In other words, either the detailed half or the general form is sufficient to tell the difference between two tokens of binoculars, for example. Although not parsimonious (two types of memory representations leading to the same result), in both the solid and stripe occluded conditions, either representation may have been detailed enough to identify which picture had been previously seen.

The remaining experiments explore these options and replicate and extend Experiment 1. In Experiments 2, 3, and 4, we examined the possibility that the memory representations in both occluded conditions were equally useful by manipulating the visual information that was available at either encoding or test. Experiments 5 and 6 use different patterns of occlusion and memory test stimuli to examine the long-term memory representation further.

Experiment 2

Although trying to remember an occluded object is a challenge for long-term visual memory, often the situation is reversed; one has to recognize that a partially occluded object is a previously seen fully visible object. For example, is that partial coffee cup my favorite coffee cup? This situation is similar to the challenges of anorthoscopic perception where one tries to identify the shape of a stimulus when it is being viewed through a narrow slit (Rock, 1996). The stimulus or the slit can move, allowing the observer to identify the shape of the stimulus, but there are limitations of the types of motion and the characteristics of the slit that will allow the shape to become apparent. Analogous to this situation, in Experiment 2, we presented only fully visible objects for encoding and had stripe or solid occluded objects, along with the fully visible control, in the memory test. Could participants match the occluded test image to the image of the complete object that they had previously stored? In this case, the encoded representations of the objects would be identical across conditions, but during the memory test, the occluded objects would be currently visible, and thus, the participant could use the available perceptual information to match to memory. If the reason that solid and stripe occlusion produced the same memory performance in Experiment 1 was because the two representations of the occluded conditions, although perceived as different, were ultimately transformed into a representation containing the same information, then in Experiment 2 we should see different memory performance when there is no need to store a representation of the occluded object.

Method

Experiment 2 used the same materials as Experiment 1 with the exception that during the encoding phase, fully visible objects were shown, and at test, objects were presented fully visible, with solid occlusion, or with stripe occlusion. Examples of how the test stimuli would be presented were shown to the participant prior to the test phase of each block. Except for that one change, the procedure was identical to Experiment 1. A new group of 72 participants were recruited and were naïve to the purpose of the experiment.

Results and Discussion

Similar to Experiment 1, memory test accuracy was affected by the occlusion condition, F (2, 138) = 15.64, $p < .001$, $\eta_p^2 = .185$, with memory in the fully visible condition being best compared to the occlusion conditions, both ts (71) > 4.00, ps < .001. Once again, however, no difference in memory performance was found between the occlusion conditions (See Table 2). With respect to the condition order, although there was no main effect, $F < 1$, there was a significant interaction between the occlusion condition and condition order, F (4, 138) = 3.38, $p = .011$, $\eta_p^2 = .089$. The interaction was primarily the result of the stripe condition that had slightly worse memory performance in the first block compared to the other blocks, whereas the other conditions were best in the first block. Similar to Experiment 1, we analyzed the first block results separately to eliminate any concern that the overall pattern was the result of carryover effects. The same pattern of memory results held in this reduced analysis with the fully visible objects (M = .88, SD = .077) being remembered better than either solid (M = .80, SD = .106) or stripe (M = .78, SD = .096) occluded objects, F (2, 71) = 8.416, p = .001. Memory in the two occlusion conditions did not differ, t (46) = .466, p >.20.

As in Experiment 1, we analyzed the effect of the location of the occluding element. Again, although there was a main effect of the location of the occluder, F (3, 213) = 11.45, p < .001, η_p^2 = .139, there was no interaction with the occlusion condition, F < 1. Objects occluded on the bottom or right were equally well remembered. However, bottom and right occluded objects were remembered better than both top and left occluded objects (Bonferroni corrected ps < .01 except the bottom to left comparison, Bonferroni corrected p = .069). Importantly, even though there were differences based on the location of the occlusion, the memory pattern was the same for both of the occlusion conditions.

The results of Experiment 2 indicate that even when the encoded representation was as complete as it can be and the memory test stimuli were occluded, memory performance was unaffected by the pattern of occlusion. In both occlusion cases, however, memory performance was substantially better when the fully visible object was encoded (Experiment 2) than when an occluded object was encoded (Experiment 1); memory was better for both the solid occlusion, t (142) = 2.84, p =.003, and stripe occlusion, t (142) = 4.02, p < .001, in Experiment 2 than Experiment 1. Because the fully visible condition in Experiments 1 and 2 were identical, it is not surprising that there was not a difference between the memory performance in the two experiments in that condition, t (142) = 1.23, p =.220.

Even given the changes from Experiment 1 to Experiment 2, the same general pattern of the effects of occlusion on memory emerged. In both experiments, it is important to note that there was a mismatch between the encoded representation and the tested stimuli except in the fully visible condition. Experiment 3 explored whether the mismatch could be limiting the memory performance in either of the occlusion conditions and whether the two conditions were differentially affected by the change in visual information from study to test.

Experiment 3

In the previous experiments, the encoded and test stimuli were different in the occlusion conditions. In Experiment 1, the encoded image was occluded, whereas in Experiment 2, the test image was occluded. Although these are common situations in the world, the mismatch between what is encoded and what is tested could place an artificial limit on memory performance. Transfer appropriate processing (e.g., Morris, Bransford, & Franks, 1977; Roediger, Weldon, & Challis, 1989; Weldon & Roediger, 1987) supports the idea that the match between the encoding and test processing can affect memory performance. Related to the concept of encoding specificity (Tulving & Thomson, 1973), transfer appropriate processing predicts that when encoding and test processing match, memory will be better than when encoding and test processing fail to match. In the two previous experiments, memory performance in the occlusion conditions may have suffered from the mismatch of the study and test stimuli. Although unlikely, the mismatch between an occluded encoded representation and a fully visible test image could be the sole reason that memory performance in the occlusion conditions was worse than in the fully visible condition (where the study and test stimuli were the same). A more likely possibility is that the mismatch between the study and test stimuli would affect one of the occlusion conditions more than the other. This mismatch of the memory representation and the test stimulus should be more limiting to memories from a condition that emphasized more detail (i.e., the solid condition) because those memory representations would be less similar to the memory test stimuli than a representation emphasizing the whole form of the object. In other words, the fact that there is a mismatch between study and test could create an artificial ceiling in memory performance that is more limiting in the solid condition. In order to evaluate the possible contribution of the match/mismatch between the study and test phases, Experiment 3 matched the study and test presentation of the objects.

Method

Experiment 3 used the encoding materials from Experiment 1 and the test materials from Experiment 2. Thus, for the occlusion conditions, objects were studied with 50% of the area occluded and were tested with the same image. For the fully visible condition, Experiment 3 replicates the previous two experiments. Like in Experiment 2, each encoding and test phase began with an example of how the objects would look in that particular block. Seventy-two participants who had not participated in Experiments 1 or 2 were recruited from the human participant pool at California State University San Marcos.

Results and Discussion

As in the other experiments, memory test accuracy was affected by encoding condition, F (2, 138) = 22.61, $p < .001$, $\eta_p^2 = .247$, and memory performance was best in the fully visible condition compared to the occlusion conditions, both ts (71) > 3.60, ps < .002 (See Table 2). In contrast to the previous experiments, memory for solid occluded objects was better than that of stripe occluded objects, t (71) = 2.496, $p = .015$. There was no main effect of condition order, F (2, 69) = 1.78, $p = .177$, $\eta_p^2 = .049$, but there was a significant interaction between the occlusion condition and condition order, F (4, 138) = 7.15, $p < .001$, $\eta_p^2 = .172$, similar to the other experiments. Unlike the previous experiments, the fully visible condition did not show much decrease in memory performance across blocks (<1%), but the two occlusion conditions did; the difference between Blocks 1 and 3 for the solid condition was 9.5% and the difference for the stripe condition was 7.8%. The lack of a block difference in the fully visible condition could be the result of the fact that fully visible objects only appeared in one block for each participant in the experiment rather than in all of the blocks (at either encoding or during the memory test) . The ability to differentiate this condition may have reduced the proactive interference that occurred across blocks in the other experiments.

Once again, we analyzed Block 1 performance as a separate between-participants ANOVA to determine if the overall pattern was the result of carryover in successive blocks. Unlike the previous experiments, there was a change in the pattern when examining only the first block. There was a marginal overall effect of condition, F (2, 69) = 2.439, p = .095, but unlike the previous results, the memory performance in the solid occlusion condition (M = .874, SD = .094) was statistically equivalent to memory in the fully visible condition (M = .865, SD = .082) in Block 1. In addition, the solid condition had better memory performance than the stripe condition (M = .823, SD = .082), t (46) = 2.014, p = .05. The fully visible condition did not differ significantly from the other conditions.

We also analyzed the effect of the location of the occluding element on the occlusion condition and found that, although there was a main effect of the location of the occluder, F (3, 213) = 4.21, p = .006, η_p^2 = .056, there was no interaction with the occlusion condition, F < 1. The main effect in location was the result of memory being worse for objects occluded on the left compared to those occluded on the bottom (Bonferroni corrected p = .022). No other *post hoc* comparison was significant.

The finding that solid occlusion resulted in better memory than stripe occlusion overall, and equal to fully visible in the first block, demonstrates that the memory representations of the two occlusion conditions are not equal. The solid condition was created to emphasize the visual information for part of the object instead of its whole object form. The visual details of the visible portion appear to be more easily encoded in visual memory, allowing the participant to more easily recognize them. However, this greater ease of encoding the visual information of part of an object does not help when one is forced to make a comparison to a completely visible object. The mismatch of visual information of a complete object and a representation of partial object resulting from occlusion may be similar to the findings that features present in whole faces are claimed to be represented differently than those same features seen on their own (Tanaka & Farah, 1993). Fully visible objects may be qualitatively different than occluded objects, making it difficult to match a representation of an occluded object to the fully visible object regardless of the form of occlusion. However, when the encoding and test match, the visual details of part of the object can be easily extracted and compared.

An interesting corollary of this argument is that solid occluded objects would appear to be represented in a way that emphasizes only the visible portion of the object (i.e., half an object). Memory for the solid condition was likely better in Experiment 3 than in the previous experiments because the representation held in long-term visual memory was the same as the visual information provided at test. In contrast, there was not as large of an advantage for matching the stripe occlusion encoding and test, indicating that how those objects were represented in long-term visual memory did not retain the specific details about the occluding element. We will explore the predictions of this corollary in Experiments 5 and 6.

Experiment 4

Although we found that matching the encoding and testing stimuli affected memory for the solid occlusion, it is unlikely that one would encounter the same object with exactly the same occlusion twice (the situation in Experiment 3). A more likely situation is that a person would see an occluded object, change his or her viewpoint, and see the same object with different parts occluded. The object remains occluded, but the change in viewpoint alters what is visible. In Experiment 4, we examined if changing the viewpoint of an object during encoding could lead to a more complete representation. By shifting the occluding mask half way through the one second encoding period, we could determine if participants are able to integrate two halves of an object into a complete whole. If participants are able to integrate the visible information of an object from two views, then the stored memory representation would be of an object with no occlusion (similar to the fully visible condition), and testing memory with a fully visible object would match the representation. This situation is analogous to perceiving an object through a slit in anorthoscopic perception (Rock, 1996) except that the slit in one case is very large (solid condition) and in the other case is divided into five separate slits (stripe condition). The difference in the demands of the occlusion conditions could make this integration more or less successful.

Method

The same stimuli and design from Experiment 1 were used in Experiment 4. The only change occurred in the procedure. Instead of presenting the object statically for one second, after 500 ms, the image changed so that the opposite occlusion image was displayed for the remaining 500 ms. For example, if the top of the object had been occluded initially, the image was replaced with an image of the object with the bottom half occluded, ensuring that the entire image was presented to the participant during the encoding period. In the stripe condition, the change shifted the phase or cycle of the stripes revealing the parts of the object that were previously occluded. In both occlusion conditions, the shifting of the location of the occluder(s) meant that the portion that was visible for the first half of the encoding time was masked for the second half of the time and vice versa. The masking should have eliminated visible persistence of the change to the stimulus (Irwin & Thomas, 2008). A change in image occurred in all of the conditions; however, in the fully visible condition, the fully visible object was replaced at 500 ms with an identical image resulting in no apparent change occurring. As in Experiment 1, the images during the memory test were fully visible. A new group of 72 participants were recruited for this experiment.

Results and Discussion

Memory test accuracy was affected by occlusion condition, F (2, 138) = 19.88, $p < .001$, η_p^2 = .244, with the same pattern as in Experiment 3, fully visible followed by solid and then stripe (See Table 2). Performances in all of the conditions were significantly different from each other, ts (71) > 2.65, ps < .02. There was no main effect of condition order, F (2, 69) = 0.16, p >.50, η_p^2 = .005, but there was a significant interaction between the occlusion condition and condition order, F (4, 138) = 5.46, p < .001, η_p^2 = .137, similar to the other experiments. As in Experiment 3, memory performance in the fully visible condition did not decrease across blocks (~2.5%) compared to the two occlusion conditions; the difference in memory between Blocks 1 and 3 for the solid condition was 6.7% and 8.1% for the stripe condition.

We analyzed Block 1 performance as a separate between-participants ANOVA, but there was no overall difference between the different conditions, F (2, 69) = 1.38, p = .26, though the condition order was in the same direction as the overall analysis: fully visible (M = .827, SD = .129), solid occlusion (M = .798, SD = .128), and stripe occlusion (M = .770, SD = .098). The lack of a difference in the first block indicates that participants were relatively able to integrate the visual information across the shift in the occluder. Even though there was no statistical difference between the different conditions in the first block, the pattern in Block 1 and the overall difference found in the complete analysis indicates that the ability to integrate across the conditions is not equal. It appears that it is easier to integrate the visual information in the solid condition where half the object is visible and there is only one occlusion edge compared to the stripe condition where integration would occur across the entire object and across multiple occlusion edge boundaries.

The encoding situation in Experiment 4 may have allowed participants to create a visually complete representation of an object in long-term visual memory similar to the fully visible condition (similar to anorthoscopic perception where a complete object is perceived through moving behind a slit). Instead of allowing the participant to shift attention around the object normally, the shifting occluder could act to serially direct attention to different portions of the object, especially in the solid condition. The serial direction of attention is not equivalent to seeing the entire object, but attending to large contiguous parts of objects appears to allow one to remember it better than a situation where one is trying to process the form of the entire object (the stripe condition). In other words, shifting the occluder in the solid condition can add visual information to a representation because new details can be revealed and added to the representation. However, when the occlusion emphasizes a representation of the whole form of the object (i.e., the stripe condition), shifting the occluder does not add as much visual information because the whole form remains the same.

Experiment 5

The previous experiments used the same occlusion conditions and varied when and how objects would be occluded. The two occlusion patterns were chosen with the assumption that the solid condition would emphasize the visual details of part of an object and the stripe condition would emphasize the overall form. Our assumption appears consistent with the results of the previous experiments, but if correct, it has additional ramifications for the representations in long-term visual memory. In order to investigate these ramifications, we ran Experiments 5 and 6.

Experiment 5 investigated if participants could remember *exactly* what they encoded of occluded objects. Our interpretation of the previous experiments was that participants could remember the portion of the object that had been occluded specifically in the solid condition, but they remember the general form of the stripe occluded objects. If correct, then when asked to remember exactly what was occluded of an object and what was visible, we should see a clear difference in the memory performance where solid occlusion results in better memory than stripe occlusion. In the memory test in Experiment 5, we presented participants with the exact occluded picture used for encoding (as in Experiment 3) with a foil that was the same object token with the complementary occlusion pattern. For example, if the red teapot was presented with the top half occluded during encoding, the same red teapot would be presented with the top occluded as the presented image and the bottom occluded as the foil image in the memory test (See Figure 2). Because the same object token was presented in both test images, the only accurate way to identify the presented image is remembering what portions of the object were visible. Because this sort of test is impossible with the fully visible condition, we introduced a new occlusion pattern that was intended to be even more disruptive than the stripe and solid condition: the checker condition (See Figure 2).[1]

Method

Seventy-two participants who had not participated in the other experiments were recruited for this experiment. Experiment 5 was based on Experiment 3 with a new occlusion condition. Checker occlusion was created by modifying the same 10 x 100 pixel bars used in the stripe and solid occlusions. Each of the five bars was divided into ten 10 x 10 pixel squares (approximately .5° x .5°), and each of these 50 squares was positioned on the object image alternating with visible 10 x 10 pixel squares, comparable to how black and white squares are arranged on a chessboard or checkerboard (see Figure 2). Again, as in the solid and stripe conditions, 50% of the image area was occluded. Because of the alternating pattern, only two occlusion patterns were possible with checker occlusion (unlike the four possible arrangements for the other conditions) that can be distinguished by whether the upper left corner is occluded or visible.

The checker condition replaced the fully visible condition in the block sequences in Experiment 5 (the solid and stripe encoding portions were identical to the previous experiments). The other change to the method of Experiment 3 was that instead of the memory test foil being a different object token, the foil was the same object token with the complementary occlusion pattern. In other words, if the binoculars were occluded on the left side during the encoding phase, that picture (left side occluded binoculars) would be tested against the same picture of binoculars but with the right side occluded. Participants were shown examples of how objects would appear

[1] A separate experiment was performed with a new group of 72 participants that tested the stripe and checker occlusion patterns with fully visible objects on the memory test (similar to Experiment 1) along with another occlusion pattern. In this experiment, memory for the stripe condition (M = .739, SD = .117) was similar to the same condition in Experiment 1 (M = .732). The checker condition (M = .687, SD =.096) led to worse memory than the stripe condition, $t(71) = 3.27$, $p = .002$, but memory for the checker occluded objects was significantly better than chance performance ($p < .001$).

before each encoding and test phase, and were informed before encoding that they would have to recognize the exact picture that they were shown. All other aspects of Experiment 5 were the same as Experiment 3.

Results and Discussion

Although the memory test was substantially more difficult in Experiment 5, memory performance in all occlusion conditions was significantly greater than chance (.50, all *t*s > 4.0, *p*s < .001). The overall effect of occlusion condition, F (2, 138) = 133.58, p < .001, η_p^2 = .659, showed that the memory performance in the solid occlusion condition was better than the other two conditions, *t*s (71) > 11.00, *p*s < .001 (See Table 3), and memory performance in the stripe condition was better than the checker condition, t (71) = 3.65, p < .001. There was no main effect of condition order, F (2, 69) = 0.97, p >.30, η_p^2 = .027, and there was only a marginal interaction between the occlusion condition and condition order, F (4, 138) = 2.42, p = .051, η_p^2 = .066. For consistency, we analyzed Block 1 performance as a separate between-participants ANOVA and found that the condition pattern was the same as the overall pattern (solid M = .779, SD = .113; stripe M = .605, SD = .093; and checker M = .542, SD = .064), F (2, 69) = 42.78, p < .001).

Finally, we analyzed the effect of the location of the occluding element on the occlusion condition for only the stripe and solid conditions to be consistent with the other experiments. We found that, although there was a main effect of the location of the occluder, F (3, 213) = 6.18, p < .001, η_p^2 = .080, there was no interaction with occlusion condition, F (3, 213) = 1.31 p = .27, η_p^2 = .018. In Experiment 5, the main effect of occlusion location resulted from the fact that memory performance in the top occlusion condition was better than memory performance in the other three locations (Bonferroni corrected *p*s < .03). No other pairwise comparison was significant. The finding that top occlusion resulted in better memory than the other conditions may indicate that, even though it may be harder to remember generally (see Experiments 1-3), because the test was to remember the location of the occluder, considerable effort went into encoding of the top occluded object initially, providing a more accurate representation of what was missing. Although it is a speculative argument, better memory for the location of an occluder while simultaneous worse memory for the object would support that possibility.

The results of Experiment 5 indicate that when tasked with remembering exactly what was seen previously, memory performance in the solid occlusion (M = .748) was similar to that in the solid condition in our previous experiments (Experiment 1 M = .738, Experiment 2 M = .787, Experiment 3 M = .817, Experiment 4 M = .765). This level of performance was more impressive because in Experiment 5, the participant could not rely on other visual details that might distinguish one token from another. The only information that could allow participants to remember the correct picture was remembering specific details about how the object was occluded. In contrast, memory performance in the stripe condition in Experiment 5 (.579) dropped precipitously compared to the previous experiments (Experiment 1 M = .732, Experiment 2 M = .794, Experiment 3 M = .786, Experiment 4 M = .728). The substantial drop in accuracy for the stripe condition in Experiment 5 indicates that the information about what could or could not be seen of the object at encoding was not precisely encoded. The inability to identify exactly what was seen originally mirrors the results of Intraub and Richardson (1989), who found that participants do not remember scenes exactly as they were shown; instead, memories for scenes contained more information that had to be inferred by the participant (see Hubbard, Hutchison, & Courtney, 2010 for a review of this literature). Some details may be remembered, accounting for above chance performance (possibly a stripe blocked a particular memorable feature of the object), but the representation generally fails to include these details. If participants are viewing the striped objects and representing the general forms, then it would be difficult to recall which portion of the object was or was not shown previously because both images in the memory test would equally match the stored representation.

Experiment 6

Because a test of exact memory appears to give an advantage to remembering the detailed side of the object (i.e., the solid condition), another test that would give an advantage to a more generalized form representation should have an advantage for the stripe and checker condition, but a disadvantage for the solid condition. The results of Experiment 5 indicate that participants cannot remember the portions of the object occluded in stripe and checker conditions, but instead appear to represent the object as a more general form. If that is the case, it creates a rather interesting prediction: for stripe and checker conditions, participants should be able to easily identify the token of the encoded object even when shown portions of the object never seen. In other words, if participants are shown the complementary occlusion stripe or checker pattern for the presented object (the 50% of the object not shown during study) during the memory test, they should be able to "remember" the opposite part of the object (see Figure 2). In contrast, because the object representation contains relatively detailed information about the part of the object that was visible, not the unseen portion, performance in the solid occlusion condition should be severely limited when tested on the complementary half.

Method

The materials from Experiment 3 were used again in Experiment 6 with two modifications. The first modification was that we used the checker condition as in Experiment 5. The more critical manipulation occurred during the memory test. As in Experiment 3, participants performed a token discrimination task, but rather than showing the presented object with its studied occlusion pattern, we presented the encoded object with the complementary occlusion pattern. By presenting the complementary occlusion pattern of the presented object as the target on the memory test (the right half of the binoculars when the left half was studied), the memory test stimulus had 0% overlap with the visible area of the studied object. Another way to think about this situation is that the foils from Experiment 5 became the targets of the memory test and were tested against a color-category matched foil. In other words, participants were performing a memory test for pictures that they had never seen. A new group of 72 participants were informed before encoding that they would have to recognize the opposite half of the presented objects, and were shown examples of images prior to each encoding and test phase.

Results and Discussion

Even though we were testing memory for images never shown, memory performance in all occlusion conditions was significantly greater than chance (.50, all *t*s > 10.0, *p*s < .001). Consistent with our prediction, memory performance in the stripe and checker conditions was better than memory in the solid condition, F (2, 138) = 17.10, p < .001, η_p^2 = .199 (See Table 3). Memory for objects in the stripe condition was significantly better than memory in the checker condition, t (71) = 2.18, p = .032, and the solid condition, t (71) = 5.17, p < .001. Interestingly, object memory in the checker condition was also better than memory in the solid condition, t (71) = 3.57, p = .001. There was no main effect of condition order, F (2, 69) = 0.43, p >.50, η_p^2 = .012, but there was a significant interaction between the occlusion condition and condition order, F (4, 138) = 6.18, p < .001, η_p^2 = .152. Once again, all conditions dropped from Block 1 to Block 3, but the stripe condition had the largest drop (8.7%) followed by the checker condition (5.8%) and the solid condition (2.4%). We analyzed Block 1 performance as a separate between-participants ANOVA as in the other experiments. The overall pattern was replicated in this reduced analysis, F (2, 69) = 5.02, p = .009, with the same relative order of conditions (stripe M = .735, SD = .119, checker M = .682, SD = .132, and solid M = .620, SD = .124), but only the stripe-solid comparison was significantly different, t (46) = 3.26, p = .002.

Consistent with Experiment 5, we analyzed the effect of the location of the occluding element on the occlusion condition for only the stripe and solid and found that, like in every other experiment, there was a main effect of the location of the occluder, F (3, 213) = 3.51, p = .016, η_p^2 = .047, and there was no interaction with the occlusion condition, F (3, 213) = 1.48 p = .22, η_p^2 = .020. The only significant difference of occluder location was that the originally encoded right occluded objects were remembered better than the originally encoded top occluded objects (Bonferroni corrected $p = .037$). It is important to remember that participants had to select the bottom occluded test image for the originally encoded top occlusion location. The worse performance in this condition (it was the worst for both stripe and solid) again provides some support to the notion that representing the top of the object is more advantageous than the other portions. However, the *post hoc* nature of this comparison does warrant caution.

As predicted, memory for objects occluded by stripes or checker patterns outpaced that for objects occluded by a solid occluder, indicating that the representations of the striped and checkered objects were likely general form representations that allowed one to gloss over the gaps caused by the occlusion. Because the representation is of the whole form of the object, it makes it difficult to identify what was precisely shown—the difference between the stripe and checker between Experiments 5 and 6 were both significant, ts (142) > 6.0, ps < .001—but allows for the participant to envision the other parts of the object and easily recognize them later. However, although the representation was useful enough to identify the non-shown parts, memory performance did not reach the level of the stripe condition in Experiment 3, where the study and test matched and the foil object was a different object token, $t(142) = 5.35$, $p < .001$. This finding indicates that although the long-term visual memory representation is able to address/“fill in” gaps, some visual details are retained that can boost memory even more.

With respect to the solid condition, the results of Experiments 5 and 6 indicate that the memory representations of the solid occluded objects contain details about the exact image that was shown, with little of the generalized form representation that can be found in the stripe or checker conditions. Memory performance for solid occluded objects was substantially better when participants had to perform a difficult exact picture decision (Experiment 5) than when asked to discriminate different unseen tokens (Experiment 6, $t(142) = -7.40$, $p < .001$). The representations of these solid occluded objects appear to contain detail about the visible parts of the object with little or no information about occluded parts stored. Because memory was still above chance in Experiment 6, however, some information about the unseen parts was able to be used. We would speculate that instead of representing the missing/occluded visual information in long-term visual memory, participants try to determine if the other side of an object could match one of the representations that is stored, like a jigsaw puzzle. Without some amount of overlap between what was stored and what is viewed on the memory test, it is very difficult to find the matching piece in long-term visual memory resulting in poor memory performance.

General Discussion

The current study explored how object representations in long-term visual memory were affected by the occlusion of some parts of the visual information of those objects. We explored two primary ways that objects could be occluded, solid or stripe, both of which should have affected the memory representations. Similar to seeing a bicycle behind a wall, where only the front part of the bicycle is visible, the solid occlusion allows the viewer to see considerable detail from the visible portion, but nothing from the occluded portion. In contrast, the stripe occlusion is similar to seeing the same bicycle behind a picket fence. The fence's slats block out discontinuous parts of the object, allowing the viewer to see and represent the overall form of the bicycle. Although the same total area is occluded in these two situations, the memory representations should be different. This difference, if present, would allow us to explore how flexibly memories are stored in long-term visual memory.

Differences in memory performance can be informative in testing what can and cannot be stored in long-term visual memory. Although the intent in using the different occlusion patterns was to force the participant to represent the information differently, the results of the first two experiments indicate that when a representation of an occluded object is tested with a non-occluded object (or vice versa), the information that is available in either pattern of occlusion was sufficient to perform the memory test. A few possible options were raised to explain the almost identical performance in the stripe and solid occlusion conditions in these experiments. One possible option was based purely on the amount of the object occluded regardless of pattern; in both cases 50% of the object area was occluded. Although that idea would account for the data in the first and second experiment, it is challenged by our other experiment results. Specifically, the findings in Experiments 3 and 4, as well as the separate study mentioned in Footnote 1, argue against this relatively low-level account. In each of those cases, there was a difference in memory performance between occlusion conditions even though the total occlusion level was the same.

The second option raised to explain the data pattern from the first two experiments was that although we intended participants to extract different memory representations based on the patterns of occlusion, they could have represented the objects in exactly the same way. Because the representation would have been the same, it would be expected that memory performance would be the same. Again, this notion was demonstrated to be incorrect because we later demonstrate that the memory representations (Experiments 5 and 6) are qualitatively different. The information stored in the solid occlusion condition allows for more precise information about the visible portions to be stored in long-term visual memory. In contrast, the stripe occlusion pattern permitted more abstract, general form representations that were not precise enough to identify exactly what was seen (Experiment 5), but were abstract enough to allow one to “fill in the gaps” as indicated by the ability of participants to “remember” the parts of the object never shown to them. Performance in the stripe condition when tested on parts never seen (Experiment 6, M =.686) approached the level of performance when the entire object was revealed by a moving occluder (Experiment 4, M =.728).

Because the first two options were not supported by the data of the other experiments, we concluded that two different types of representations could be created and that the memory test in Experiments 1 and 2 was insensitive to that difference. In other words, one could use either the abstract form or the details that were visible equally effectively when moving from the occluded objects to the fully visible objects or vice versa. This explanation has the disadvantage of claiming that two separate ways of representing an object will lead to the same performance, but it is consistent with the results of the other experiments. In Experiments 3-6, we tested whether the representations from the different occlusion patterns could result in different memory performance. We found that in testing memory for exactly what was shown (Experiments 3 and 5), there was an advantage to representing a detailed side of the object over the disconnected stripes of the entire object. However, when the test favored the use of the general form of the object (Experiment 6), we clearly saw an advantage for the disconnected, but whole object, representations.

What does this study mean for long-term memory representations overall? First, long-term visual memory is flexible, allowing multiple representation forms depending on the information available during encoding. Memories are not “converted” or “transformed” into a universal format that contains the same information every time. Instead, memories can retain substantial visual detail, general form information, or both depending on the encoding situation. Neither specific details nor general form information is qualitatively better, and they can frequently lead to similar memory performance (depending on the test). If, however, the memory representation contains both form information and details of the object, it will lead to better memory performance (in all cases, the fully visible condition was better than either occlusion condition).

Second, when presented with the bicycle-behind-a-wall situation described above, people may believe that an entire bicycle exists, but in representing the object in visual memory, they don’t appear to attempt to create a memory of the complete object (including the entire side that was not visible), instead opting for a detailed memory of the visible portion. Even when participants were aware of the fact that they would be tested on the occluded portion of the object, they performed poorly in the solid occlusion condition. In contrast, when the visible portions of an object are separated by a gap, like in the picket fence-bicycle example, participants created a memory that generalized across the whole form of the object (similar to simple geon compatible completions, Biederman, 1987). This generalized abstract form allows one to “remember” what was not seen (Experiment 6), but not remember well the exact image seen (Experiments 3 and 5). The representation has a more generalized form that does not exactly match up to the visual details of the tested object (i.e., transfer appropriate processing, Morris et al., 1977).

Third, although the form is generalized in the picket fence situation, some specific details must be maintained in the representation because participants were substantially above chance performance in identifying the exact picture studied (see also Henderson, 1997, for a similar finding in transsacadic memory). The only way to remember which stripe pattern was presented is to remember that a specific detail (e.g., the seat of the bicycle) was visible during the encoding of the object. Likely, the ability to represent a specific detail will be idiosyncratic to both the object and the viewer. If a particular visible detail is interesting or special, the representation of the object will contain that information. However, the default procedure for representing objects occluded in a discontinuous manner will emphasize form over these details.

How occluded objects are represented in long-term visual memory speaks to the flexibility and limitations of the long-term visual memories, but what of the connection between the sorts of amodally completed representations typically discussed in the attention literature and those studied here? Previous work on amodal completion (e.g., Palmer et al., 2006) has emphasized that people fill in information behind the occluder. Given that amodal completion can occur during online object processing, why do we find that completion sometimes appears to be present and sometimes absent? We would argue that the difference lies in the requirements of the task. In the previous research, amodal completion was required to perform the task (e.g., shape judgement; Murray et al., 2001). One could not perform the task unless the occluded object could be completed. In contrast, in the current study, the occluding information was merely an additional difficulty added to the task, not the requirement. Remembering a bicycle does not require trying to imagine the missing components of the bicycle; one can identify the category of the object without the unseen components and represent what is visible (see Konkle et al., 2010, about the importance of conceptual category in memory). However, when the occlusion pattern encourages processing of the object form in order to identify the category of the object, one can use an abstract form that can "fill in" the missing information. This difference in task between the previous research and the current study would seem to be a prime reason why we find that some representations appear to contain more precise information whereas others are more abstract. The amodal completion of long-term memory representation is not obligatory, but rather is dependent on what will be necessary to categorize the object.

The flexibility of long-term visual memory representations permits the system to use whatever information is perceptually available to encode and remember what was seen. Although one can find differences between representation types, under most viewing situations, the ability to use either a more precise detailed representation of a partial object or the overall form of an object allows the individual to remember the object fairly well. In the current study, the encoding conditions encouraged one type of memory representation or another because that was what was episodically available during encoding. Extending this idea, if a viewing task itself would emphasize one type of representational format or another, we would argue that the memory representation of even fully visible objects would be biased toward containing the emphasized information. Although we referred to these as "detailed" and "form" representations, we do not think that long-term visual memory is limited in that manner. Anything that is emphasized or relevant at the time should affect long-term visual memories. Antonelli and Williams (2016) have shown that memory for search targets is different when a specific perceptual attribute is part of the target definition (blue backpacks) than when the same targets are defined without that attribute (backpacks). The results of Antonelli and Williams and the current study both point to importance of the encoding context for long-term visual memory. Further exploration of the influence of task demands on visual memory will allow a more complete understanding of the episodic nature of these memories.

Summary: Overall, long term visual memories appear to have multiple forms depending on the relevant information that is available at encoding. The representation can be highly detailed about a portion of the object or a more abstract, generalized form. Each of these types of representations may be sufficient to perform one type of memory task, but they can fail when trying to access the information in a different manner (i.e., transfer appropriate processing). Although there is a lack of parsimony in the number of representations that can be created, the combination of the available stimulus information and the memory task permits a system that represents what is critical at the moment to the viewer, and thus may create a more useful long term visual memory representation.

References

Antonelli, K. B., & Williams, C. C. (2016). Episodically relevant features define categories in visual memory. Manuscript submitted for publication.

Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, *94*(2), 115-147. doi:10.1037/0033-295X.94.2.115

Brady, T. F., Konkle, T., Alvarez, G. A., & Oliva, A. (2008). Visual long-term memory has a massive storage capacity for object details. *PNAS Proceedings of the National Academy of Sciences of the United States of America, 105*(38), 14325-14329. doi:10.1073/pnas.0803390105

Castelhano, M. S., & Henderson J. M. (2005). Incidental memory for objects in scenes. *Visual Cognition, 12*, 1017-1040. doi:10.1080/13506280444000634

Henderson, J. M. (1997). Transsaccadic memory and integration during real-world object perception. *Psychological Science, 8*(1), 51-55. doi:10.1111/j.1467-9280.1997.tb00543.x

Henderson, J. M., & Hollingworth, A. (2003). Global transaccadic change blindness during scene perception. *Psychological Science, 14*(5), 493-497. doi:10.1111/1467-9280.02459

Hollingworth, A., & Franconeri, S. L. (2009). Object correspondence across brief occlusion is established on the basis of both spatiotemporal and surface feature cues. *Cognition, 113*(2), 150-166. doi:10.1016/j.cognition.2009.08.004

Hollingworth, A., & Henderson, J. M. (2002). Accurate visual memory for previously attended objects in natural scenes. *Journal of Experimental Psychology: Human Perception and Performance, 28*(1), 113-136. doi:10.1037/0096-1523.28.1.113

Hubbard, T. L., Hutchison, J. L., & Courtney, J. R. (2010). Boundary extension: Findings and theories. *The Quarterly Journal of Experimental Psychology, 63*(8), 1467-1494. doi:10.1080/17470210903511236

Intraub, H., & Richardson, M. (1989). Wide-angle memories of close-up scenes. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15*(2), 179-187. doi:10.1037/0278-7393.15.2.179

Irwin, D. E., & Thomas, L. E. (2008). Visual sensory memory. In S. J. Luck & A. Hollingworth (Eds.), *Visual memory* (pp. 9-41). New York, NY, USA: Oxford University Press, Inc.

Kellman, P. J., & Shipley, T. F. (1992). Perceiving objects across gaps in space and time. *Current Directions in Psychological Science, 1*, 193-199. doi:10.1111/1467-8721.ep10770407

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

Morris, C. D., Bransford, J. D., & Franks, J. J. (1977). Levels of processing versus transfer appropriate processing. *Journal of Verbal Learning & Verbal Behavior, 16*(5), 519-533. doi:10.1016/S0022-5371(77)80016-9

Murray, R. F., Sekuler, A. B., & Bennett, P. J. (2001). Time course of amodal completion revealed by a shape discrimination task. *Psychonomic Bulletin & Review, 8*, 713-720. doi:10.3758/BF03196208

Palmer, E. M., Kellman, P. J., & Shipley, T. F. (2006). A Theory of dynamic occluded and illusory object perception. *Journal of Experimental Psychology: General, 135*, 513-541. doi:10.1037/0096-3445.135.4.513

Rayner, K. (2009). Eye movements and attention in reading, scene perception, and visual search. *The Quarterly Journal of Experimental Psychology, 62*(8), 1457-1506. doi:10.1080/17470210902816461

Rensink, R. A., & Enns, J. T. (1998). Early completion of occluded objects. Vision Research, 38, 2489-2505. doi:10.1016/S0042-6989(98)00051-0

Roediger, H. I., Weldon, M. S., & Challis, B. H. (1989). Explaining dissociations between implicit and explicit measures of retention: A processing account. In H. I. Roediger & F. M. Craik (Eds.), *Varieties of memory and consciousness: Essays in honour of Endel Tulving* (pp. 3-41). Hillsdale, NJ, US: Lawrence Erlbaum Associates, Inc.

Rock, I. (1996). Anorthoscopic perception. In I. Rock (Ed.), *Indirect perception* (pp. 107-124). Cambridge, MA: MIT Press.

Schneider, W., Eschman, A., & Zuccolotto, A. (2002). E-Prime (Version 1.2) [Computer Software]. Pittsburgh, PA: Psychology Software Tools, Inc.

Sekuler, A. B., & Palmer, S. E. (1992). Perception of partly occluded objects: A microgenetic analysis. *Journal of Experimental Psychology: General, 121*, 95-111. doi:10.1037/0096-3445.121.1.95

Standing, L. (1973). Learning 10,000 pictures. *The Quarterly Journal of Experimental Psychology, 25*, 207-222. doi:10.1080/14640747308400340

Tanaka, J. W., & Farah, M. J. (1993). Parts and wholes in face recognition. *The Quarterly Journal of Experimental Psychology A: Human Experimental Psychology, 46A*(2), 225-245. doi:10.1080/14640749308401045

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(6), 1494-1505. doi:10.1037/0096-1523.21.6.1494

Tulving, E., & Thomson, D. M. (1973). Encoding specificity and retrieval processes in episodic memory. *Psychological Review, 80*(5), 352-373. doi:10.1037/h0020071

Weldon, M. S., & Roediger, H. L. (1987). Altering retrieval demands reverses the picture superiority effect. *Memory & Cognition, 15*(4), 269-280. doi:10.3758/BF03197030

Williams, C. C. (2010a). Incidental and intentional visual memory: What memories are and are not affected by encoding task? *Visual Cognition, 18*, 1348-1367. doi:10.1080/13506285.2010.486280

Williams, C. C. (2010b). Not all visual memories are created equal. *Visual Cognition, 18*, 201-228. doi:10.1080/13506280802664482

Williams, C. C., Henderson, J. M., & Zacks, R. T. (2005). Incidental visual memory for targets and distractors in visual search. *Perception & Psychophysics, 67*, 816-827. doi: 10.3758/BF03193535

Wolfe, J. M., Reijnen, E., Horowitz, T. S., Pedersini, R., Pinto, Y., & Hulleman, J. (2011). How does our search engine 'see' the world? The case of amodal completion. *Attention, Perception, & Psychophysics, 73*(4), 1054-1064. doi:10.3758/s13414-011-0103-0

Table 1.

Experimental condition for the six experiments in the study.

Experiment	Encoding Conditions	Memory Test Conditions
Experiment 1	Fully Visible, Solid, Stripe	Fully Visible
Experiment 2	Fully Visible	Fully Visible, Solid, Stripe
Experiment 3	Fully Visible, Solid, Stripe	Fully Visible, Solid, Stripe
Experiment 4	Fully Visible, Solid, Stripe	Fully Visible
	(Occluder moved after 500 ms)	
Experiment 5	Solid, Stripe, Checker	Solid, Stripe, Checker
		(Exact portion occluded tested)
Experiment 6	Solid, Stripe, Checker	Solid, Stripe, Checker
		(Opposite occluded portion tested)

Table 2.

Mean memory results (SD in parentheses) for Experiments 1-4. For Experiments 1 and 4, Condition refers to how the objects were presented at encoding with fully visible memory test objects. For Experiment 2, Condition refers to how the objects were presented during the memory test with objects encoded fully visibly. For Experiment 3, because the encoding and memory test presentation of the objects was the same, Condition refers to both the encoding and memory test presentation of the objects.

	Condition		
Experiment	Fully Visible	Solid	Stripe
Experiment 1	.830 (.111)	.738 (.107)	.732 (.093)
Experiment 2	.852 (.104)	.787 (.101)	.794 (.094)
Experiment 3	.863 (.074)	.817 (.110)	.786 (.106)
Experiment 4	.811 (.118)	.765 (.117)	.728 (.103)

Table 3.

Mean memory results (SD in parentheses) for Experiments 5 and 6. In both Experiments 5 and 6, Condition refers to how the objects were presented at encoding and during the memory test.

	Condition		
Experiment	Solid	Stripe	Checker
Experiment 5	.748 (.116)	.579 (.088)	.533 (.068)
Experiment 6	.616 (.098)	.686 (.117)	.656 (.119)

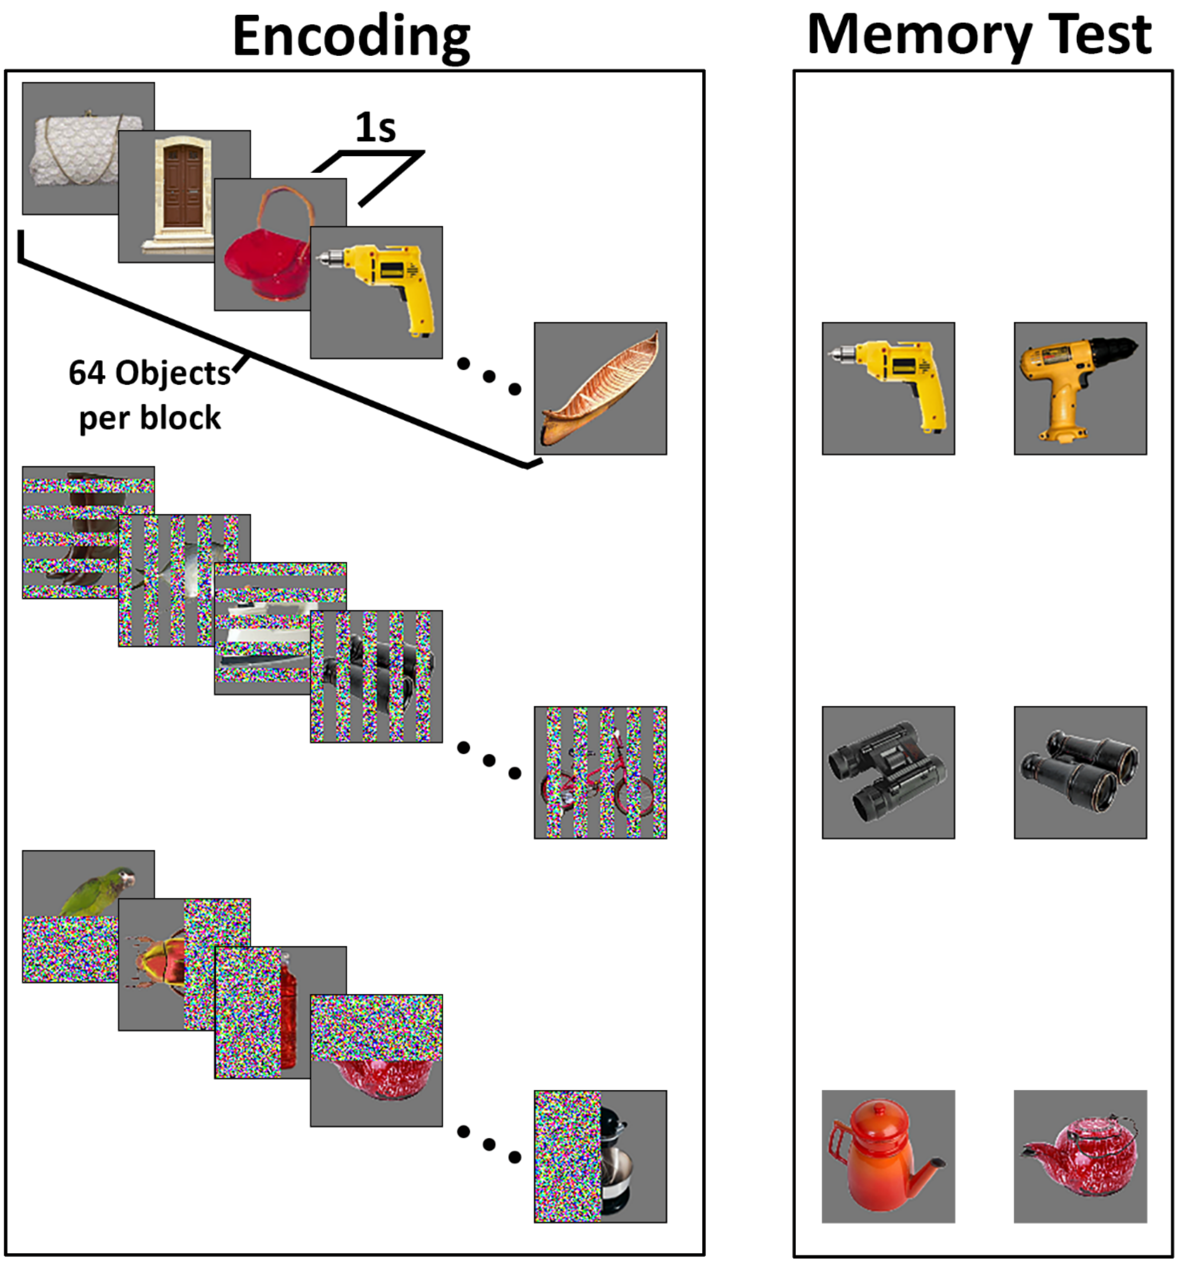

Figure 1. A description of the procedure and stimuli from Experiment 1. The order of block was counterbalanced across participants. Stimuli were presented in full color on a neutral gray background (RGB=120). Within an encoding block of 64 objects, all pictures were presented with the same condition (Fully visible, Stripe, or Solid). Memory tests followed the encoding sequence in each block.

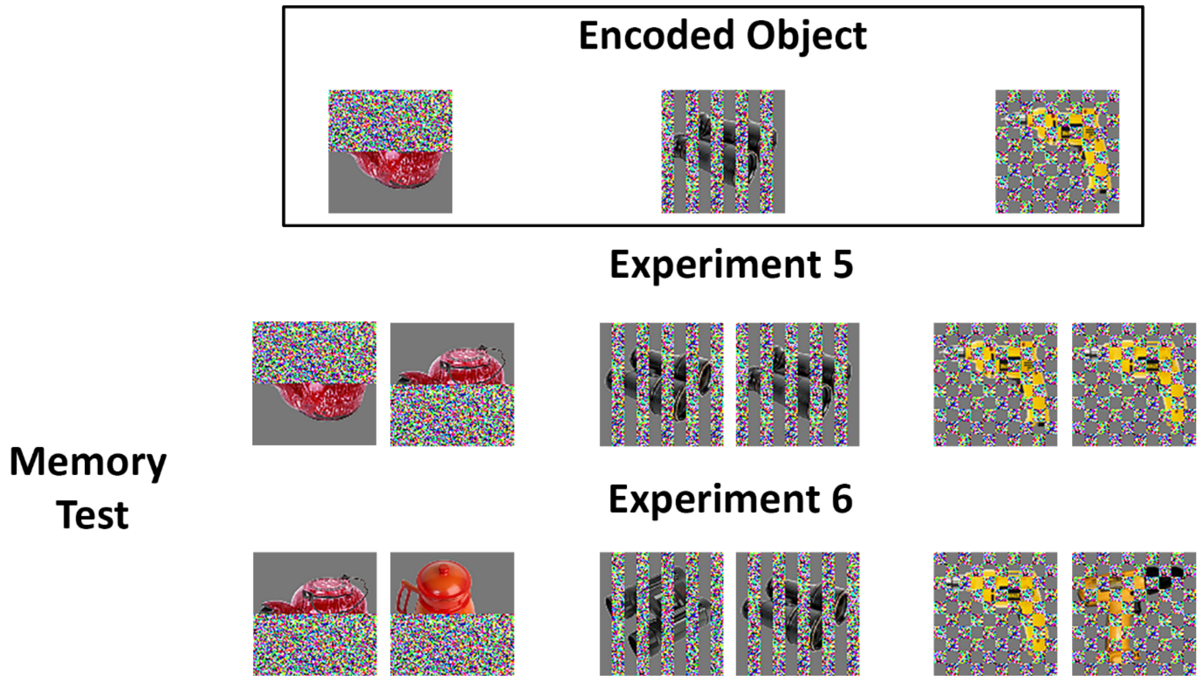

Figure 2. Example of the stimuli and test conditions for Experiments 5 and 6. If during the encoding phase, participants saw (in separate blocks) the solid occluded teapot, a stripe occluded binoculars, and a checker occluded drill (as shown at the top of the figure), the memory tests would appear like those below. In Experiment 5, participants were instructed to choose the image that was identical to the one encoded (the exact image). In Experiment 6, participants were instructed to choose the object token encoded even though the visible parts of the test object were the previously occluded portions (i.e., the non-presented portions of the object). The correct answers in both examples are: Solid = Left, Stripe = Right, Checker = Left. Images were presented in color.