Investigating Dissociations Between Perceptual Categorization and Explicit Memory

Marci A. Flanery (marci.flanery@vanderbilt.edu)
Thomas J. Palmeri (thomas.j.palmeri@vanderbilt.edu)
Brooke L. Schaper (brooke.l.schaper@vanderbilt.edu)
Department of Psychology; Vanderbilt University
Nashville, TN 37240 USA

Abstract

Are dissociations between categorization and explicit memory in tests of amnesics and normals evidence for multiple memory systems? Or could these dissociations be artifacts arising from methodologies used in some experiments? We report a series of studies exploring this issue. Using normals in various states of simulated amnesia we show that categorization at test is well above chance even in the absence of prior exposure to category members. We also show that subjects perform well when tested with items that conflict with categories they had studied earlier. We argue that subjects in some paradigms can extract information about categories from the test rather than rely on memory for studied category members. In further studies, we generalize these findings to other stimuli and other category structures that have been used in tests of amnesics and normals.

Introduction

Do categorization and explicit memory rely on independent neural memory systems? Evidence for multiple systems comes from dissociations between categorization and explicit memory in studies of normals and amnesics. Amnesics are reported to categorize at levels comparable to normals but are significantly worse at explicit memory. Such dissociations seem to imply that separate systems may exist and pose clear problems for theories that assume a single underlying memory system, such as well-known exemplar models.

The evidence is clear that amnesics have impaired explicit or declarative memory. The focus of this paper is on whether data from studies testing amnesics clearly provide evidence for intact abilities to learn new perceptual categories. Our goal is to examine whether some categorization performance can be explained in the absence of positing a separate implicit system for category learning that is spared in amnesia. Our approach has been to utilize the same paradigms and methodologies found in the amnesia literature to study normal subjects under conditions that simulate aspects of amnesia. To create "amnesia" in normals, we used a variety of techniques such as eliminating the study session altogether, introducing delays between study and test, and surreptitiously switching the test stimuli to those from an unstudied category. We follow the amnesia literature in testing these effects using a variety of stimuli, including distortions of dot patterns, object-like stimuli with discrete features, and simple forms placed in categories separated by quadratic boundaries.

In this paper, we review some behavioral evidence for dissociations from studies of amnesics and normals. For each case, we present data from studies we conducted that provide a possible alternative explanation for intact categorization by amnesics. Due to space constraints, we will only present our results in summarized form without detailed description of the methods or statistical analyses. After summarizing our initial work along these lines reported by Palmeri and Flanery (1999), we describe several new experiments that expanded upon these initial results in several important ways.

Learning Categories of Dot Patterns

A classic methodology for studying categorization and recognition has been the Posner and Keele (1968) dot pattern paradigm. To create a pattern, a small number of dots are randomly scattered on a grid. To create a category, a pattern is randomly generated and designated the prototype. Category members are generated by randomly distorting the prototype by varying degrees.

Knowlton and Squire (1993) used a variant of this paradigm to test amnesics on categorization and recognition. For categorization, subjects were exposed to 40 high distortions. Subjects were tested on judging members and nonmembers of that category. For categorization, members were 4 repetitions of the prototype, 20 low distortions, and 20 high distortions. Nonmembers were 40 randomly generated patterns. For recognition, subjects were exposed to five random patterns eight times each. In the recognition test, they were asked to discriminate between the five old patterns and five new patterns. No corrective feedback was provided in either condition. Knowlton and Squire (1993) reported a dissociation between categorization and recognition when comparing amnesics and normals. As shown in Figure 1, amnesics categorized as well as normals but were significantly impaired at recognition memory.

This dissociation seemed to provide evidence for separate systems. However, Nosofsky and Zaki (1998) showed that a single-system model could account for a

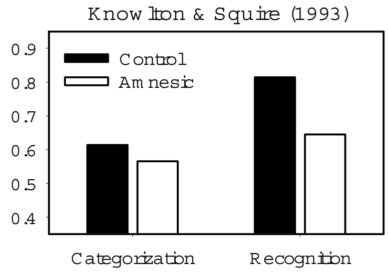

Figure 1. Categorization and recognition accuracy for controls and amnesics.

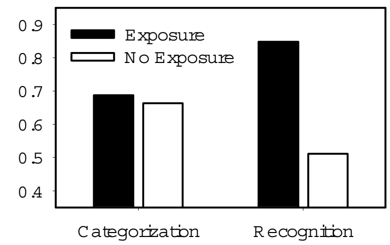

Figure 2. Categorization and recognition accuracy in Experiment 1 for subjects exposed and not exposed to category items (Palmeri & Flanery, 2001).

dissociation by simply assuming that amnesics had degraded memory. A challenge for this theoretical possibility was an extreme dissociation observed by Squire and Knowlton (1995). Their patient, E.P., was able to categorize as well as normals but recognition was entirely at chance. It would be very difficult to formulate a single-system model along the lines of Nosofsky and Zaki that could predict chance recognition performance in the presence of normal categorization performance.

To better understand the categorization performance of amnesics, Palmeri and Flanery (1999) investigated whether prior exposure was even necessary to categorize at test. One explanation for above-chance categorization by amnesics is that it may be possible to group items during the test that looked similar (prototypes and distortions) into the member category and group items that did not look similar (random patterns) into the nonmember category. Whereas, it is impossible to tell apart old from new patterns without memory.

Palmeri and Flanery tested this possibility by producing a state of profound amnesia in normals. As a ruse, subjects were told that patterns had been subliminally presented during an initial word identification task. No dot patterns were ever really presented. Subjects then completed the same categorization and recognition tests used by Knowlton and Squire. Similar to E.P., our simulated profound amnesics showed chance recognition. Yet, they showed above chance categorization. Apparently, our subjects were able to figure out how to categorize members versus nonmembers by picking up on the category structure clearly embedded within the test. They had no prior memories of any sort to rely on.

Experiment 1. We extended this paradigm by directly comparing the performance of simulated amnesics (No Exposure) to that of subjects who were exposed to the study items (Exposure). Half of the subjects were given subliminal exposure, as in Palmeri and Flanery, and were tested on categorization or recognition; the other half were given actual exposure, as in Knowlton and Squire, and tested on categorization or recognition. Results are shown in Figure 2. As expected, the exposure group could recognize items well above chance but the no exposure group could only guess. Replicating Palmeri and Flanery (1999), subjects in a no exposure group could categorize well above chance. Subjects receiving no exposure did not categorize significantly worse than subjects who were actually exposed to category items. Apparently, prior exposure to a category did not provide much, if any, benefit.

Experiment 2. One criticism of these studies is that the ruse used to induce amnesia may have placed subjects in a different mind set from that of subjects who were exposed to members. Our "profound amnesics" may realize they never saw any patterns and may think the task is to discover the hidden category structure, something they appear to do quite ably. So, one goal of this experiment was to use a different paradigm for demonstrating that subjects may categorize based on information they acquire during the categorization test. In this experiment, we surreptitiously switched the test stimuli for some subjects to that of an unstudied category.

In addition, we clearly do not want to draw the conclusion that people always ignore information about a previously studied category in favor of information presented during a test. A second goal was to show that when initial exposure provides evidence for a clear category structure, subjects will use that information to make category decisions irrespective of the makeup of the categorization test. To demonstrate this, we adapted a paradigm used by Squire and Knowlton (1995). In one condition, subjects were initially exposed to 40 high distortions of the prototype (40H), exactly as was done in earlier studies. In another condition, subjects were instead exposed to 40 repetitions of the prototype (40P). We reasoned that subjects in the 40P condition should have acquired clear knowledge of the category structure and should protest any surreptitious changes during a test. By contrast, subjects in the 40H condition should have acquired little knowledge of the category structure and should go along with our surreptitious changes.

To verify that different exposure conditions had a significant effect on performance, we tested subjects in the same way as our earlier studies after a one week delay. Overall, 40P subjects achieved 75.3% accuracy and 40H subjects achieved 65.1% accuracy. As expected, categorization accuracy was influenced by the type of information presented during initial category exposure, as was reported by Squire and Knowlton (1995). Overall performance of our 40H subjects was quite comparable to what we and others have observed in this paradigm; performance of the 40P subjects was significantly better than what we have observed before. So, information presented during initial exposure can have a significant effect on categorization performance.

As a way of simulating amnesia, we tested a subset of subjects after a several weeks delay. But now we tested half on items generated from the prototype used to generate items they had seen before (Same condition) and tested half on items generated from a novel prototype (Different condition). Thus, each subject was in one of four conditions: 40P-Same, 40P-Different, 40H-Same, and 40H-Different. Since all subjects viewed a different randomly generated set of stimuli, we can characterize subjects in the Different condition as receiving a categorization test intended for another individual.

As illustrated in Figure 3, we found that subjects in 40P-Same performed quite well, correctly categorizing 77.4% of the items. However, subjects in 40P-Different were completely at chance categorizing the test items. We suspect that these subjects tried to use the category information they clearly had acquired earlier and could not apply that knowledge when given a test comprised of entirely novel items. For subjects in the 40H conditions, as we predicted, there was no significant difference in performance between subjects who were tested on the same structure they were initially exposed to and subjects who were tested on a completely different structure. Consistent with our previous results, these subjects appear to be making categorization decisions based on what they acquired during the categorization test, not on what they may have acquired during earlier phases of the experiment.

Summary. The dissociation between categorization and recognition reported by Knowlton and Squire (1993) initially appeared to present strong evidence supporting multiple memory systems theory. We reported how the observed dissociations between categorization and recognition using distorted dot patterns may be explained as a result of the particular methodologies used to test these individuals. We showed that very good categorization performance can be achieved in the absence of any prior exposure to the category members. We also showed that very good categorization performance can be achieved when people are tested on items that are

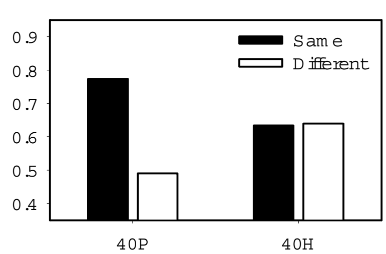

Figure 3. Categorization accuracy in Same condition and Different condition for subjects studying 40 repetitions of a prototype (40P) and 40 high distortions of prototype (40H) in Experiment 2.

different from what they had actually studied. But this seems to only occur when subjects have been initially exposed to a very diffuse category structure consisting of high distortions that are not very similar to one another. When subjects have been exposed to a clear category structure through repetition of a single prototype, they attempt to categorize items based on that acquired category knowledge, not on information presented during the categorization test.

Learning Categories of Object-Like Stimuli with Discrete Features

Reed et al. (1999) aimed to generalize the investigation of preserved categorization by amnesics by using object-like stimuli with discrete features. The stimuli they used, called Peggles, were drawings of animals that varied on nine binary-valued dimensions. To create a category, some Peggle was designated the prototype. Category members were distortions of the prototype. Low distortions shared 7 or 8 features of the prototype. High distortions shared only 1 or 2 features. In the extreme, an antiprototype had all nine features opposite to that of the prototype. Stimuli that shared 4 or 5 features of the prototype were designated neutral items that were half way between the prototype and the antiprototype.

Subjects viewed 40 low distortions of the prototype. After exposure, subjects were told that the animals they saw were all members of a category, called Peggles, and were then asked to make member/nonmember judgments, without feedback, of 96 new items. The test included 12 repetitions of the prototype, 24 low distortions, 24 neutral items, 24 high distortions, and 12 repetitions of the antiprototype. Subjects were also tested on their ability to complete a cued-recall test identifying both values of the 9 dimensions of the Peggles.

Reed et al. (1999) found that amnesics were impaired at an explicit cued-recall task but could categorize at levels comparable to normals. But, two of their amnesics actually categorized stimuli opposite to the way they should have. That is, they mistakenly called the prototype and low distortions nonmembers and called the antiprototype and high distortions members. Reed et al. suggested that amnesics had a spared implicit category learning system that partitioned members and nonmembers but that perhaps declarative memory was needed to remember which partition corresponded to the items they had previously been exposed to.

Experiment 3. Following the theme of this paper, we propose an alternative explanation. During the categorization test, subjects were shown the prototype many times and were shown low distortions that were very similar to the prototype. They were also shown the antiprototype many times and were shown high distortions that were very similar to the antiprototype. In other words, there were two clear clusters of items presented during the categorization test. If subjects could pick up on the category structure embedded within the testing sequence to cluster stimuli into two groups, they would be able to correctly partition the stimuli into two categories. But, they would not be able to unambiguously decide which cluster corresponded to the category they were initially exposed to without relying on memory of some sort. Might this be a more reasonable explanation of the category switching by amnesics previously reported by Reed et al.?

The goal of this experiment was to test whether subjects might be categorizing in part by extracting information from the structure of the categorization test. We tested subjects in three conditions: Immediate, Delayed, and Novel. The Immediate condition was essentially a replication of Reed et al. (1999). In the Delayed condition, subjects were exposed to the category and then returned one week later to be tested in the same way as subjects in the Immediate condition. In the Novel condition, subjects were also exposed to the category and returned one week later. The stimuli presented for categorization in the Novel condition contained an embedded category structure that contradicted what was presented during initial exposure. To do this, a neutral item with respect to the prototype that was used to generate stimuli from the original exposure session was picked at random and designated the "prototype" for purposes of creating a new categorization test sequence. From this novel prototype, low distortions, neutral items, high distortions, and an antiprototype were created. Note that the "antiprototype" for this new structure would also be considered a neutral item with respect to the prototype that was used to generate items subjects were originally exposed to. The novel categorization test consisted of 12 repetitions of the novel prototype, 24 low distortions, 24 neutral items, 24 high distortions, and 12 repetitions of the novel antiprototype.

Let us generate some predictions for the Novel condition. If subjects are categorizing based on what they had been previously exposed to, they should categorize the "prototype" and the "antiprototype" in this novel test sequence equally, as halfway between the member and nonmember category with respect to what they had originally studied. However, if subjects are instead picking up on the clear category structure embedded within this novel test sequence, they should group the "prototype" and its distortions in one category and group the "antiprototype" and its distortions in another category. Half of the subjects would call the "prototype" group the members and half would call the "antiprototype" group the members.

Now to the results. First, as shown in the right portion of Figure 4, performance in the recall task was significantly impaired in the Delayed and Novel condition compared to the Immediate condition. Also, as shown in the left of Figure 4, subjects in the Immediate and Delayed conditions showed comparable categorization.

Scoring categorization data for subjects in the Novel condition was somewhat more complicated (Palmeri & Flanery, 2001). Essentially, what we first did was to

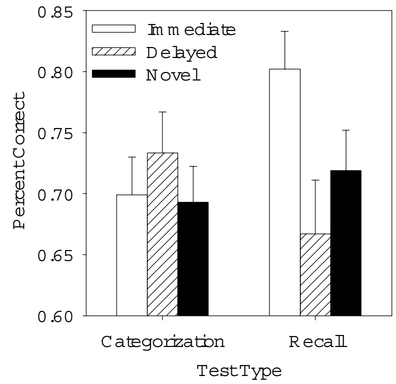

Figure 4. Categorization and cued-recall accuracy in the Immediate, Delayed, and Novel conditions of Experiment 3 (Palmeri & Flanery, 2001).

measure the difference in membership endorsements for the "prototype" and the "antiprototype." Recall that if subjects were categorizing these two critical stimuli with respect to what they had actually been exposed to, they should be indifferent at categorizing these items as members or nonmembers. To the contrary, we found a 53.6% difference in membership endorsements for the "prototypes" and the "antiprototypes." Subjects were clearly discriminating between these items when making category member judgments. Next, if a particular subject judged the "prototype" more often to be a member then we judged categorizations of the low distortions as members and high distortions as nonmembers to be "correct" responses; if a particular subject judged the "antiprototype" more often to be a member then we judged categorizations of the high distortions as members and low distortions as nonmembers to be "correct" responses. Figure 4 displays categorization accuracy for the Novel condition using this scoring method (actually, we scored the Immediate and Delayed conditions in the same way to make the reported results consistent). What should be clear from the figure is that subjects in the Novel condition discriminated between members and nonmembers in a way that was consistent with the structure embedded within the testing sequence and not on memory for what they had seen a week earlier. As with Experiment 2, we found comparable performance between subjects who were tested on categories they actually studied and subjects who were tested on categories that contradicted what they had actually studied.

Summary. In this experiment, we extended a paradigm used by Reed et al. (1999) to contrast categorization and recall by amnesics and normals. They observed impairments in cued recall by amnesics compared to normals, but there was little difference in categorization between the two groups. However, they did observe that two of their amnesic individuals categorized members of the previously studied category as nonmembers and nonmembers as members. While Reed et al. interpreted these results in terms of an implicit memory for the category, we instead provided evidence that this ability to discriminate members from nonmembers might emerge from a clear category distinction embedded within the testing sequence.

Learning Categories Described by a Complex Quadratic Rule

Can individuals with explicit memory impairments learn to categorize stimuli in accordance with a complex categorization rule? Filoteo et al. (in press) had normals and amnesics learn categories described by what they characterized as a complex quadratic rule. Subjects learned two categories that were defined by multivariate normal distributions. Figure 5 displays the equal likelihood contours for the category structures utilized by Filoteo et al. (in press). Because the categories are defined by normal distributions, the two categories overlap, so perfect performance is impossible. Also as shown in the figure, learning the categories required subjects to integrate information across both stimulus dimensions; in the language of decision boundary theory, learning these categories required the formation of a quadratic (nonlinear) decision rule. This manipulation was of theoretical importance because some work has suggested that amnesics cannot integrate information across multiple dimensions (Rickard & Grafman, 1998).

The stimuli used by Filoteo et al. (in press) consisted of a horizontal and a vertical line connected at the top left corner. The length of the horizontal and vertical lines varied in accordance with the distributions shown in Figure 5. It is important to note that the "diagonal" distribution consisted of stimuli for which the line lengths are highly correlated; in other words, they form the left and top portions of a square (square category). The "circular" distribution consisted of stimuli for which the line lengths are uncorrelated; in other words, they form the left and top portions of various rectangles (rectangle category). On each trial of the experiment, subjects were presented with a stimulus randomly drawn from either the square or the rectangle category, categorized it as a member of category A or category B, and received corrective feedback.

Filoteo et al. observed the accuracy in the last 100 trials to be 85% for normals and 84% for amnesics. They concluded that amnesics appear to be able to acquire categories defined by a complex quadratic rule. To test whether an amnesic could retain that rule over a delay period, they tested one amnesic and one normal after a one day delay. Subjects completed a single block of 100 trials in which they received corrective feedback on every trial, just as in the original training. Accuracy was 92% for the normal individual and 89% for the amnesic. Thus, amnesics and normals appear to be able to learn and retain a quadratic categorization rule.

Experiment 4. The Filoteo et al. results suggest that amnesics can learn and retain a category described by a complex quadratic rule that requires integrating information from two stimulus dimensions, height and width. However, these stimuli can also be described in an alternative way by rotating the dimensions by 45 degrees. That is, we can alternatively describe the dimensions as shape and size. The square and rectangle categories contain stimuli of the same shape and can be categorized by a very simple shape rule rather than a complex quadratic rule. Filoteo et al. rejected this possibility, arguing that their subjects were learning a complex quadratic rule requiring an integration of information along two independent stimulus dimensions. But, we are puzzled by how these subjects were able to learn a complex categorization rule so quickly, reaching asymptotic performance after less than 100 trials, when other categorization experiments examining quadratic boundaries have required many days of training to reach asymptote.

To illustrate that subjects may not be learning a complex quadratic rule, but may instead may be learning a simple shape rule, we replicated and extended the Filoteo et al. study using three conditions. In the first condition, we used the same stimuli and category structures as Filoteo et al. (Square/Rectangle condition). In the second condition, subjects were trained on similar stimuli, but both multivariate category distributions were shifted along dimension 1. In this way, the diagonal category distribution still had height and width correlated, but their values were not equal – in other words, the stimuli were rectangles of the same shape that varied in size (Rectangle/Rectangle condition). In the third condition, we used very different stimulus dimensions of circle size and angle of a diameter line (Circle-Line/Circle-Line condition) that cannot be integrated like the height and width of line segments; these dimensions were roughly equated for discriminability with the height and width dimensions.

Performance in the Squares and Rectangles conditions were comparable (81% and 78% accuracy, respectively). Performance in the Circle-Line condition was far worse (58% accuracy). These results suggest that amnesics may not have been learning a complex quadratic categorization rule at all, but may have instead been learning a very simple shape rule.

Another issue with the Filoteo et al. (in press) results regards the retention of the categorization rule after a

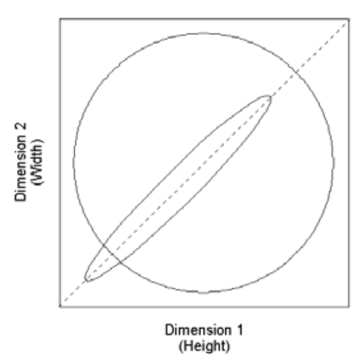

Figure 5. Equal probability contours for categories used by Filoteo et al. and used in Experiment 4.

delay. In the second session of their experiment, subjects received feedback after each trial, similar to what they had experienced during training. Did the amnesics display a real memory for the categorization rule or did they express a savings in relearning a very simple categorization rule? To show that different kinds of categorization tests can reveal different levels of knowledge about categories, we brought our subjects back after one day and tested them in three different ways. First, we tested them without feedback on stimuli drawn from a uniform distribution across the entire set of possible stimuli. Second, we tested them without feedback on stimuli drawn randomly from the two category distributions. Third, we retrained them with feedback, as was done by Filoteo et al. Although subjects reached comparable levels of performance in the Squares and Rectangles condition on the first day, subjects were much better when tested on the uniform distribution without feedback in the Squares condition (96%) than the Rectangles condition (79%). By contrast, in the other two testing conditions (without feedback and with feedback), performance was comparable for the Squares and Rectangles condition (82% and 80% accuracy, respectively). It appears that the different categorization tests can reveal differential knowledge of the categories.

Summary and Conclusions

We found evidence that normal subjects can acquire information about categories in the absence of prior study and in opposition to prior study. In our experiments, performance by subjects in these conditions was not significantly different from performance by subjects who actually received prior study and who were tested on items consistent with their prior study. Our results demonstrate that classification decisions made during a categorization test may not be based solely on information acquired during a study task, but may also be based on information acquired during the test itself. As a general point, we argue that care must be taken in selecting items for a categorization test so as not to provide additional information about the categories being tested or so as not to change the information about the categories that may have been previously acquired. In the present experiments, subjects were tested in such a way that it was possible to extract information about the categories from the tests themselves. Subjects were repeatedly tested on the category prototype (four times in Experiment 1, twelve times in Experiment 2) and were tested on many low distortions that were very similar to the category prototype, conditions particularly amendable to unsupervised category learning. A preferable way to test individuals in an neutral manner might be to sample all possible test stimuli from a uniform distribution, as we did in the last experiment. Although it may be possible to partition such test stimuli into some arbitrary set of categories, only by chance might this partition match the correct category discrimination without relying on memory for studied items.

As we stated at the outset, our results may have implications for understanding the relationship between categorization and other forms of memory. The ability of amnesics to categorize stimuli coupled with their impairment at recognizing or recalling stimuli has been taken as evidence for multiple memory systems (e.g., Knowlton & Squire, 1993; Reed et al., 1999; Squire & Knowlton, 1995; see, however, Nosofsky & Zaki, 1998). If the paradigms used by some investigators permit category acquisition from the categorization test (by contrast, the explicit memory tasks used in these experiments cannot be accurately performed without memory for the studied items), then the strength of this dissociation may be questioned. It seems prudent to forgo strong conclusions about independence or nonindependence of fundamental aspects of human cognition until more convincing paradigms are employed.

Acknowledgments

This work was supported by Vanderbilt University Research Council Grants, Grant BCS-9910756 from the National Science Foundation, and Grant MH61370 from the National Institute of Mental Health.

References

Filoteo, J.V., Maddox, W.T., & Davis, J.D. (in press). Quantitative modeling of category learning in amnesic patients. Journal of the International Neuropsychological Society.

Knowlton, B.J., & Squire, L.R. (1993). The learning of categories: Parallel brain systems for item memory and category knowledge. Science, 262, 1747-1749.

Nosofsky, R.M., & Zaki, S.R. (1998). Dissociations between categorization and recognition in amnesic and normal individuals: An exemplar-based interpretation. Psychological Science, 9, 247-255.

Palmeri, T.J., & Flanery, M.A. (1999). Learning about categories in the absence of training: Profound amnesia and the relationship between perceptual categorization and recognition memory. Psychological Science, 10, 526-530.

Palmeri, T.J., & Flanery, M.A. (2001). Category knowledge acquired during categorization testing.

Posner, M.I., & Keele, S.W. (1968). On the genesis of abstract ideas. Journal of Experimental Psychology, 77, 353-363.

Reed, J.M., Squire, L.R., Patalano, A.L., Smith, E.E., & Jonides, J. (1999) Learning about categories that are defined by object-like stimuli despite impaired declarative memory. Behavioral Neuroscience, 113, 411-419.

Rickard, T.C., & Grafman, J. (1998). Losing their configural mind: Amnesic patients fail on transverse patterning. Journal of Cognition Neuroscience, 10, 509-524.

Squire, L.R., & Knowlton, B.J. (1995). Learning about categories in the absence of memory. PNAS, 92, 12470-12474.