General Article

DISSOCIATIONS BETWEEN CATEGORIZATION AND RECOGNITION IN AMNESIC AND NORMAL INDIVIDUALS: An Exemplar-Based Interpretation

By Robert M. Nosofsky and Safa R. Zaki

Indiana University

In recent work, the finding of dissociations between categorization and recognition in amnesic and normal individuals has been taken as evidence of multiple memory systems mediating these tasks. The present research provides support for the alternative idea that these dissociations can be interpreted in terms of a single-system exemplar-memory model that makes allowance for parameter differences across groups. In one experiment, a parameter change in memory sensitivity was induced by testing classification and recognition at varying delays; the results closely matched the ones observed by Knowlton and Squire (1993) for normal and amnesic participants. The exemplar model also yielded good quantitative predictions of the categorization-recognition dissociation. A second analysis demonstrated that dissociations between early versus late probabilistic classification learning and memory sensitivity were also well predicted by the single-system exemplar model. Limitations of the exemplar interpretation and future research directions are also discussed.

According to exemplar models of classification, people represent categories by storing individual exemplars in memory, and classify objects based on their similarity to these stored exemplars. Exemplar models have had a great deal of success in predicting classification performance, and have compared favorably with alternatives such as prototype, feature-frequency, and simple rule-based models (for a review, see Nosofsky, 1992).

An important research theme motivated by exemplar models involves the study of relations between categorization and old/new recognition performance. Presumably, if individual exemplars are stored in memory during category learning, then that fact should be corroborated in postacquisition recognition memory tests. Indeed, in early research, single-system exemplar models were taken to task on grounds of certain dissociations between categorization and recognition performance. For example, researchers reported studies that demonstrated low correlations or lack of positive contingencies between classification and recognition, and argued that these results posed problems for single-system exemplar-memory models (e.g., Anderson, Kline, & Beasley, 1979, p. 314; Metcalfe & Fisher, 1986, p. 164; Omohundro, 1981).

Nosofsky (1988), however, showed that each of the dissociations was in fact well predicted by a single-system exemplar model. The key idea was that although classification and recognition may be mediated by a single representational system involving stored exemplars, the decision rules that operate in the tasks may differ (see also Estes, 1986b; Gillund & Shiffrin, 1984; Hintzman, 1986, 1988; Medin, 1986). Consider a task in which an observer learns to classify objects into multiple categories. According to the categorization decision rule, the probability that an object is classified into a target category is based on the object's similarity to the exemplars of the target category relative to its similarity to the objects of the contrast categories. By contrast, recognition is based on the absolute summed similarity of the object to all exemplars of all the categories. This absolute summed similarity gives a measure of familiarity, with higher familiarity values leading to higher recognition probabilities. If different decision rules are involved, dissociations between categorization and recognition could exist, even though similarity comparisons to stored exemplars are involved in both tasks. Beyond demonstrating that the dissociations were consistent with the exemplar model's predictions, Nosofsky (1988, 1991; Shin & Nosofsky, 1992) demonstrated that the model provided good, simultaneous quantitative fits to the fine-grained classification and recognition probabilities associated with individual objects in the stimulus sets used. Thus, the ability of the exemplar model to simultaneously predict quantitative details of classification and recognition performance provided strong converging evidence in its favor.

Address correspondence to Robert M. Nosofsky, Department of Psychology, Indiana University, Bloomington, IN 47405; e-mail: nosofsky@indiana.edu.

Recently, however, a new set of intriguing dissociations between categorization and recognition, of a different character than the earlier ones, has been reported by Knowlton, Squire, and their colleagues (Knowlton, Mangels, & Squire, 1996; Knowlton & Squire, 1993; Knowlton, Squire, & Gluck, 1994; Squire & Knowlton, 1995). These dissociations provide a new challenge to the idea that a single-system exemplar model may mediate categorization and recognition. The purpose of the present article is to begin to examine exemplar-based accounts of these categorization-recognition dissociations (for closely related work, see D'Arcy & Larochelle, 1997). Although we do not provide a complete account of these dissociations in terms of the exemplar model, we take some important first steps.

STUDY 1

Knowlton and Squire (1993) reported experiments in which groups of amnesic and matched normal control subjects categorized or made old/new recognition judgments for sets of visual patterns. In the categorization task, the stimuli were dot patterns generated from a prototype by using the classic statistical-distortion methods of Posner and Keele (1968). Participants viewed 40 high distortions of the prototype. Following this viewing phase, participants were tested with 84 new transfer patterns including the prototype (four presentations), 20 low distortions of the prototype, 20 high distortions, and 40 random patterns. Participants judged whether each new pattern belonged to the same category as the training patterns. The participants were also tested in an old/new recognition task. During study, five random patterns were presented eight times each. During test, the five old patterns and five new random patterns were presented, and participants judged whether each one was old or new.

Knowlton and Squire's (1993) main results are summarized in Figure 1 (top panel). As expected, the normal subjects performed significantly better than did the amnesic subjects on the recognition test. The critical result, however, was that the two groups performed with similar accuracy on the categorization test (the differences were not significant). Furthermore, there was a significant interaction between the performance of the two groups on the classification and recognition tests.

A more detailed breakdown of the categorization results is shown in Figure 2 (top-left panel), which illustrates a classic *prototypicality effect.* Both groups endorsed the prototypes with the highest probability, followed in order by the low distortions, high distortions, and random patterns.

Knowlton and Squire interpreted their results in terms of multiple memory systems. In their view, an implicit memory system is responsible for the acquisition of category-level knowledge, which is represented in the form of prototypes. However, a declarative system involving memories for exemplars mediates old/new recognition. Because the declarative system in amnesics is damaged, their old/new recognition is

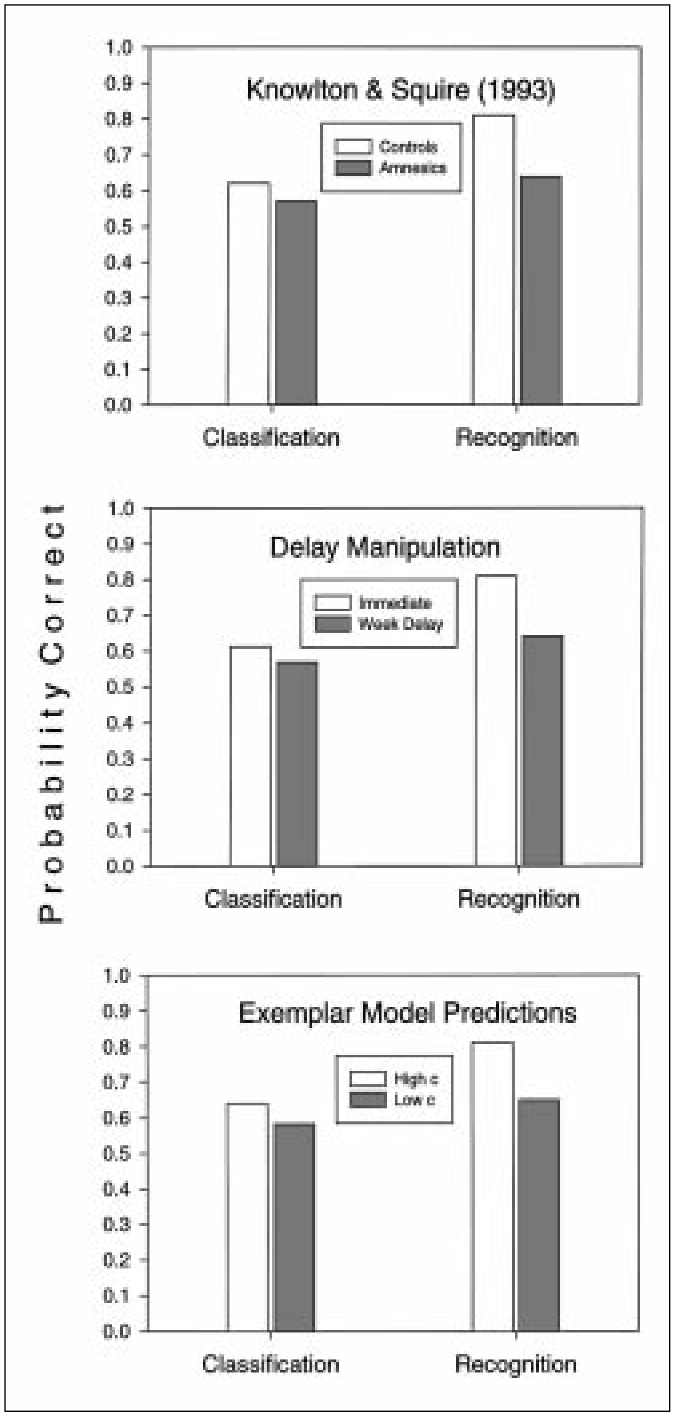

Fig. 1. Observed and predicted proportions of correct classification and recognition decisions. The top panel shows the proportion of correct classification and recognition decisions made by the normal and amnesic participants in Knowlton and Squire's (1993) experiment. The middle panel shows the proportion of correct classification and recognition decisions made in the immediate and 1-week-delay conditions of Study 1. Predictions from the exemplar model with high and low settings of the memory-sensitivity parameter (c) are graphed in the bottom panel.

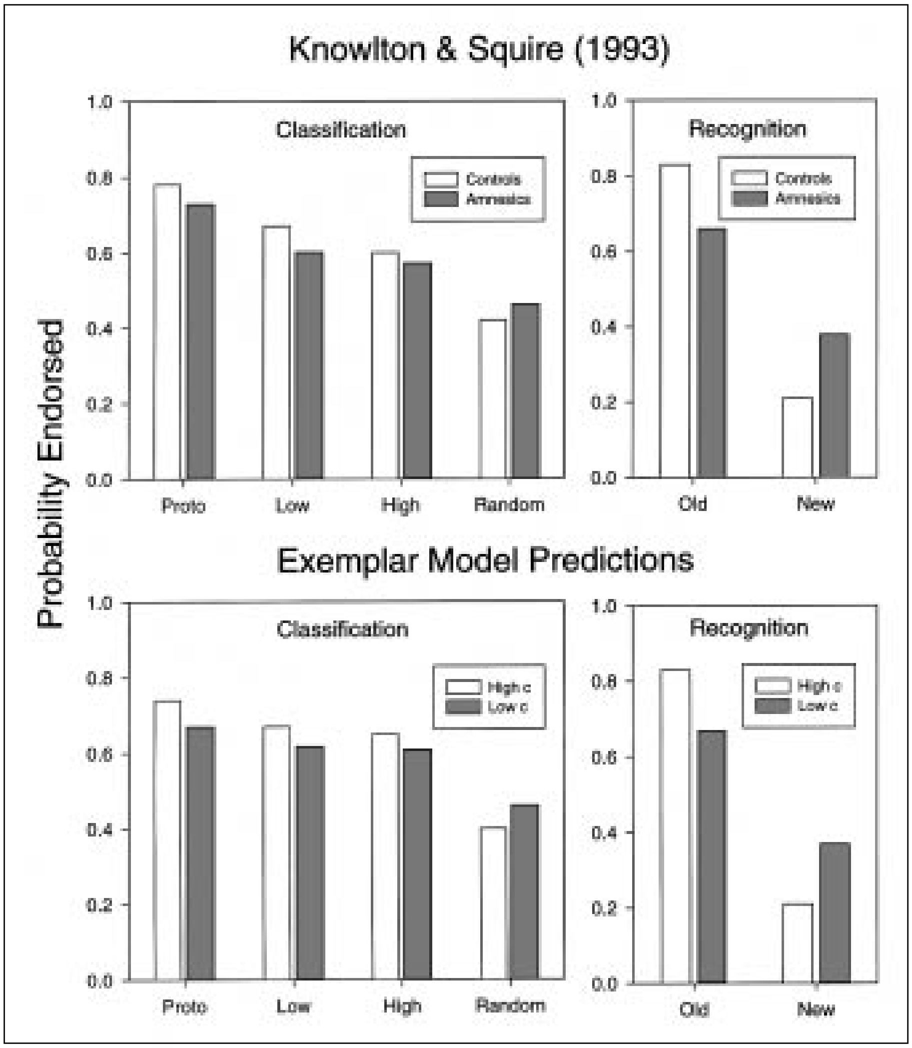

Fig. 2. Observed and predicted proportions of classification and recognition decisions for each type of pattern. The top panels show the probability with which each type of pattern was endorsed as a category member or recognized as old in Knowlton and Squire's (1993) experiment. Exemplar-model predictions of Knowlton and Squire's (1993) data are graphed in the bottom panels. Proto = prototype; Low = low distortions of the prototype; High = high distortions of the prototype; c = memory-sensitivity parameter.

impaired. However, the implicit system in amnesics is intact, so their categorization performance does not suffer.

Furthermore, note that the categorization-recognition dissociation reported by Knowlton and Squire (1993) cannot be explained in terms of the use of a relative-similarity rule for categorization and absolute-summed-similarity rule for recognition (Nosofsky, 1988) because in Knowlton and Squire's task, only a single category was involved.

It may be possible, however, to explain Knowlton and Squire's results with a single-system exemplar model in which parameter differences exist between the two groups. Suppose that there is a single memory system and that categorization is based on similarities to stored exemplars. Presumably, the exemplar-based memory system of the amnesics is impaired. For example, the amnesics may have more difficulty than the normal individuals in discriminating among distinct exemplars in memory. This ability to discriminate is represented in terms of an overall *sensitivity* parameter in the exemplar model (Nosofsky, 1987). Perhaps the amnesics' lower level of sensitivity hurt their performance in the recognition task, but was still sufficient to support normal performance in the categorization task (Shanks & St. John, 1994).

This possibility involving parameter differences between groups was directly acknowledged by Knowlton and Squire (1993, p. 1748), and, as we discuss later, addressed in other experimental paradigms (Knowlton et al., 1996; Squire & Knowlton, 1995). Nevertheless, Knowlton and Squire's (1993) original results clearly leave the burden on exemplar theorists to demonstrate the plausibility of the parameter-difference explanation. Our first main goal in the present study was to test the extent to which a single-system exemplar-memory model, which makes allowance for parameter differences between groups, could account for the classification-recognition dissociation reported by Knowlton and Squire.

We took a two-pronged approach to meeting this goal. First, using Knowlton and Squire's (1993) stimulus materials, we tested normal participants in both the classification and recognition tasks. Our idea was to induce a parameter change across different groups by manipulating the delay of the test phase. Or, to put it another way, we sought to "simulate" amnesia by introducing delays. Presumably, as delay before testing increased, participants' memories for the exemplars would get worse, as would be revealed by significantly poorer old/new recognition. The key question was, what effect would the parameter change have on performance in the corresponding classification test?

Second, we collected similarity ratings among representative patterns used in both tasks. The goal was to use these similarity ratings in conjunction with a formal version of the exemplar model to quantitatively fit Knowlton and Squire's data. The issue was whether a single-system exemplar-memory model, which makes allowance for a parameter change across groups, could quantitatively account for the classification-recognition dissociation.

Method

Participants

The participants were 120 undergraduates at Indiana University who received partial course credit. Twenty students participated in each of six experimental conditions.

Stimuli

The stimuli were the same dot patterns used by Knowlton and Squire (1993).[1]

Procedure

Participants were trained on either classification or recognition patterns and were given a transfer test either immediately after training, 1 day later, or 1 week later. (We tested both 1-day and 1-week delays because we could not know in advance how much delay would be needed to yield significant forgetting in the recognition task.) Classification training followed Knowlton and Squire's (1993) procedures. The participants were shown the series of 40 study patterns and, for each pattern, were asked to point to the central dot on the computer screen. Each of the 40 patterns was presented once and remained on the screen for 5 s. Recognition training used the same procedure except each of the 5 study patterns was presented eight times. Order of presentation of the items was randomized for each participant.

In the classification transfer task, the participants were told that they had seen dot patterns that belonged to a category, just as if they had seen a series of pictures of the category "dog." On each trial, the task was to decide whether the presented pattern was a member of the category. In the recognition transfer task, participants judged whether or not each pattern had been presented during training. To ensure high motivation, we told the participants the top six performers in the experiment would receive a $15 bonus.

Immediately after the transfer test, the participants made similarity judgments to pairs of patterns. For each participant, a different randomly selected subset of the patterns was used (because including all patterns would have resulted in an excessive number of ratings). The items in each subset included the prototype, three low distortions, three high distortions from the training set, three high distortions from the transfer set, and three random patterns. Participants rated the similarities by using a 9-point scale (1 = least similar, 9 = most similar).

Results

The overall percentage of correct recognition and classification decisions in each of the delay conditions is reported in Table 1. In the 1-day-delay condition, there was little reduction in either recognition or classification (compared with the immediate condition), so we do not consider this condition further. The most interesting results were obtained in the 1-week-delay condition. Compared with the immediate condition, there was a large reduction in overall recognition performance. Critically, however, there was extremely little reduction in overall classification. The results of the immediate and 1-week-delay conditions are displayed in Figure 1 (middle panel). The degree of correspondence between our data and those of Knowlton and Squire (1993) is remarkable. Recognition performance was sig-

Table 1. *Proportion of correct classification and recognition decisions in each delay condition*

	Task	
Condition	Recognition	Classification
Immediate	.81	.61
1-day delay	.78	.59
1-week delay	.64	.57

1. Knowlton and Squire (1993) conducted the classification task twice using two different sets of dot patterns. We used only the first set in this study.

nificantly worse in the 1-week-delay condition than in the immediate condition, $t(38) = 3.37$, $p < .005$, but classification performance did not differ significantly across conditions, $t(38) = 1.12$, $p > .10$. There was also a significant interaction between condition and task, $F(1, 76) = 3.97$, $p = .05$, reflecting the greater drop in recognition than in classification as a function of delay.

In our view, these results increase the plausibility of the idea that the categorization-recognition dissociation reported by Knowlton and Squire may indeed be attributable to a parameter change across groups, although such an interpretation is, of course, not forced. For example, perhaps there are two separate memory systems that mediate recognition and categorization, but the recognition system suffers more rapid forgetting than the categorization system. Nevertheless, we believe that the onus begins to shift back to the multiple-memory-system account.

To provide further evidence bearing on the parameter-change explanation, we fitted a formal version of the exemplar model to Knowlton and Squire's data. According to the model, classification decisions are made by summing the similarity of an item to the 40 high distortions stored in memory and then comparing this summed similarity with a categorization criterion. Specifically, to derive formal predictions, we assumed that the probability that item i is endorsed as a member of the category is given by

$$P(\mathrm{C}) = [40*s(i,h)]/[40*s(i,h) + k_{\mathrm{C}}],$$

where $s(i,h)$ denotes the average similarity of the item to each high distortion, and k_{C} denotes the classification criterion.

Analogously, in the recognition task, the probability that an item is judged "old" is found by summing its similarity to the five study items and comparing this summed similarity with a recognition criterion. For old items, the probability of a recognition response is given by

$$P_{\mathrm{old}}(\mathrm{R}) = [\delta + 4*s(i,r)]/[\delta + 4*s(i,r) + k_{\mathrm{R}}],$$

where $s(i,r)$ denotes the average similarity between an item and the random study patterns, δ denotes the self-similarity between an item and its memory trace, and k_{R} denotes the recognition criterion. The probability of a recognition response for new items is given by

$$P_{\mathrm{new}}(\mathrm{R}) = [5*s(i,r)]/[5*s(i,r) + k_{\mathrm{R}}].$$

To use these equations, we needed parameter estimates for the interitem similarities and for the criterion settings in each condition. We used the participants' ratings to impose constraints on the estimates of the similarity parameters. The relevant data are given in Table 2, which provides the mean similarity ratings of the prototype, low distortions, new high distortions, and random patterns to the old high distortions. The table also provides the mean similarity rating among random

Table 2. *Mean similarity ratings [rating(*i,j*)] among the types of patterns*

	Pattern *j*	
Pattern *i*	Old high distortions	Random patterns
Prototype	5.311	—
Low distortions	4.961	—
New high distortions	4.863	—
Random patterns	4.004	4.139

patterns (which is relevant for computing the recognition probabilities). The similarity ratings are highly systematic. In general, the greater the distance of an object from the prototype, the lower is its average similarity to the old high distortions.

For a simple approximation, we assumed that the "true" psychological similarity between types of patterns was given by a power transformation of their rated similarity,

$$s(i,j) = [\mathrm{rating}(i,j)]^p.$$

To model the parameter differences between groups, separate power exponents (p_{N} and p_{A}) were estimated for the normal and amnesic individuals. Given our assumption of a power-model relation between "true" and rated similarity, it can be shown that a higher level of sensitivity in normal individuals translates into a higher value of their estimated power exponent (i.e., $p_{\mathrm{N}} > p_{\mathrm{A}}$).[2]

The data to be predicted are those illustrated in the top panels of Figure 2, namely, the classification and recognition probabilities yielded by the normal and amnesic participants for each individual type of item (a total of 12 response probabilities).[3] Fitting the *full* version of the exemplar model to these data required estimation of seven free parameters: the power exponents p_{N} and p_{A}, a self-similarity rating parameter δ^* [with $\delta = (\delta^*)^p$], and four criterion parameters (categorization criteria for the normal and amnesic participants, and recognition criteria for the normal and amnesic participants). Parameters that

2. According to the exemplar model, similarity (s) between a pair of objects is an exponential decay function of their psychological distance (d), $s = \exp(-c \cdot d)$, where c is a sensitivity parameter (Nosofsky, 1984; Shepard, 1987). We assume that normal and amnesic individuals exhibit different levels of sensitivity. For normal participants, set c arbitrarily equal to 1, and for amnesics, let c be some constant less than 1. Then for normal participants, $s_{\mathrm{N}} = \exp(-d)$; and for amnesics, $s_{\mathrm{A}} = \exp(-c \cdot d) = [\exp(-d)]^c = s_{\mathrm{N}}^{\ c}$. Thus, similarity between exemplars in amnesics is related by a power function to similarity between exemplars in normal individuals. Given our assumption of a power-model relation between true and rated similarity, the assertion made in the text then follows directly.

3. Knowlton and Squire (1993) reported only the overall accuracy scores for recognition, not the separate hit and false alarm rates. Because our overall recognition accuracy scores were virtually identical to theirs, we used our own obtained hit and false alarm rates as the data to be fitted in this analysis.

minimized the sum of squared deviations (SSD) between predicted and observed data were estimated.

The predicted classification and recognition probabilities for the individual item types are shown in Figure 2 (bottom panels), and the overall predicted levels of correct classification and recognition performance are shown in Figure 1 (bottom panel). The quantitative fit is quite good (SSD = 0.010), and the model captures the key trends in the data. First, as illustrated in Figure 1, the model reproduces almost exactly the overall levels of categorization and recognition performance exhibited by the normal and amnesic participants in Knowlton and Squire's (1993) experiment. Second, as illustrated in Figure 2, the model reproduces the prototypicality effect observed for both the normal and the amnesic participants, although it slightly underestimates the accuracy for the prototypes.[4] Finally, the model pinpoints the observed differences in performance between the normal and amnesic participants for each of the individual types of patterns.

The best fitting parameters are reported in Table 3. The most important result is that the value of p_N exceeds p_A, as it should if memory sensitivity for the normal participants exceeds that for the amnesic participants. The estimated value of the self-similarity parameter ($\delta = 9.72$) slightly exceeds the 9-point ceiling on the scale that participants used for rating similarities among distinct patterns. We find this result to be intuitively sensible as well. Finally, we emphasize that reasonable fits to the classification and recognition data can also be achieved by using restricted versions of the exemplar model with fewer free parameters. For example, we fitted a version in which all four criterion parameters were held fixed at values that would produce unbiased responding (equal hit and correct-rejection rates). This parsimonious three-parameter model yielded an SSD of 0.021, nearly as good as that of the full version, and produced the same pattern of results.

Table 3. *Best fitting exemplar-model parameters*

Parameter	Value
p_N	5.18
p_A	2.99
δ^*	9.72
$k_C(N)$	78,788
$k_C(A)$	2,927
$k_R(N)$	28,877
$k_R(A)$	583

Note. N = normal participants; A = amnesic participants; p = power exponent; δ^* = self-similarity; k_C = classification criterion; k_R = recognition criterion. Note that $\delta = (\delta^*)^p$.

In summary, our formal analyses indicate that a single-system exemplar-memory model that allows for a parameter change to represent the differential memory sensitivities of normal and amnesic individuals is capable of reproducing the classification and recognition probabilities that Knowlton and Squire (1993) observed for these groups.

STUDY 2

A second paradigm that yielded a dissociation between categorization and recognition in normal and amnesic participants was reported by Knowlton et al. (1994, 1996). Participants were tested in a probabilistic classification learning task. The stimuli were composed of four binary-valued dimensions that were associated probabilistically with two categories (A and B). On each trial, a stimulus was selected at random in accord with its assigned dimension-value probabilities, the participant classified it, and corrective feedback was then provided. Following classification learning, the researchers measured participants' memory for aspects of the classification task by means of a memory questionnaire.

Knowlton et al. (1994, 1996) plotted averaged learning curves in the probabilistic classification task for a group of amnesics and a group of matched normal control subjects. A key result was that during the first 50 trials of learning, the two groups did not differ significantly in their classification performance. With extended learning, however, the normal subjects performed significantly better than the amnesics. Also, the normal subjects displayed significantly better performance than the amnesics on the memory questionnaire.

Knowlton et al. (1996) suggested that early probabilistic classification learning is governed by an implicit memory system that gradually learns habits. Because this implicit system is intact in amnesics, their early classification learning does not suffer. Knowlton et al. (1996, p. 1401) suggested further that because the normal subjects outperformed the amnesics during the later trials of learning, they may eventually have made use of additional information from a declarative memory system.

In this section, we examine the extent to which the performance dissociations reported by Knowlton et al. (1994, 1996) can be predicted by an exemplar model that assumes parameter differences between groups, without the need to posit separate memory systems for different tasks and for early versus late learning. According to the model, the probability that an item is classified into Category A during any trial of the learning sequence is given by

$$P(\mathrm{A}|i) = [\Sigma s(i,\mathrm{a}) + \beta]^\gamma / \{[\Sigma s(i,\mathrm{a}) + \beta]^\gamma + [\Sigma s(i,\mathrm{b}) + \beta]^\gamma\},$$

4. One difference between the results obtained in our delay experiment and the results from Knowlton and Squire (1993) is that we observed a much smaller prototype enhancement effect, although the overall levels of performance averaged across all patterns were virtually identical across experiments. The smaller prototype enhancement effect was obtained in our immediate condition as well as in our delay conditions, so it does not seem to imply a failure of our delay manipulation to simulate amnesia. If results are averaged across experiments, the exemplar model predicts almost perfectly the magnitude of the prototype enhancement effect.

where $\Sigma s(i,\text{a})$ denotes the summed similarity of item i to all previously presented exemplars of Category A, β is a *background-noise* constant, and γ is a response-scaling parameter (Estes, 1994; McKinley & Nosofsky, 1995; Nosofsky, Kruschke, & McKinley, 1992). Early in learning, before many exemplars have been experienced, the background-noise constant dominates and responding is close to chance. As learning proceeds, the summed similarities grow in magnitude, and responding comes to be governed by the experienced exemplars. The response-scaling parameter γ describes the extent to which observers use probabilistic versus deterministic response strategies (Maddox & Ashby, 1993; McKinley & Nosofsky, 1995). When $\gamma = 1$, observers respond by using a *probability-matching* rule, whereas when γ grows greater than 1, observers respond more deterministically with the category that yields the larger summed similarity.

The methods for computing the interitem similarities [$s(i,\text{a})$] in the exemplar model have been discussed in numerous previous articles (e.g., Estes, 1986a; Medin & Schaffer, 1978; Nosofsky, 1984). In the present baseline application, we assume for simplicity that the psychological distance D between item i and an exemplar from Category A is given by the number of dimensions on which they have mismatching values. The similarity between these objects is then an exponential decay function of their distance,

$$s(i,\text{a}) = \exp(-c \bullet D),$$

where c is a scaling parameter (Shepard, 1987). The parameter c represents, in part, a participant's level of memory sensitivity in discriminating among distinct exemplars. Again, we assume that memory sensitivity is greater for the normal than for the amnesic individuals (an assumption that derives support from the results of the memory questionnaire used by Knowlton et al.).

To derive the formal predictions from the exemplar model, we constructed 1,000 random stimulus sequences that satisfied the constraints of the design used by Knowlton et al. (1994, 1996). We then used the model to predict the percentage of correct classifications in each block of 10 trials for each individual sequence, and finally averaged across the sequences to predict the averaged learning curves. In the model investigation that we report, we set $\beta = 2$ and $\gamma = 2$. We set the memory-sensitivity parameter for the amnesic group at $c = 1$, and the memory-sensitivity parameter for the normal group at $c = 3$.

The predictions from the exemplar model, illustrated in Figure 3 by the solid curves, are remarkably similar to the observed findings of Knowlton et al. (1994, 1996). During the first 50 trials of learning, there is virtually no difference between the predicted performances of the low-sensitivity and high-sensitivity groups. However, the high-sensitivity group achieves a clear advantage later in learning. The same general pattern holds for a wide range of parameter settings in the model.

The reason for these predictions can be described intuitively as follows. When sensitivity is high, there is little psychological similarity among distinct exemplars, and little in the way of stimulus generalization. The summed similarities grow slowly relative to the background memory noise, so performance moves only slowly away from chance. Once sufficient exemplar information is stored, however, it tends to be highly veridical, so performance ultimately reaches high levels. By contrast, when sensitivity is low, there is a great deal of stimulus generalization, and the summed similarities grow quickly relative to the background noise. However, because exemplars from contrasting categories are poorly discriminated, the asymptotic performance levels are low.

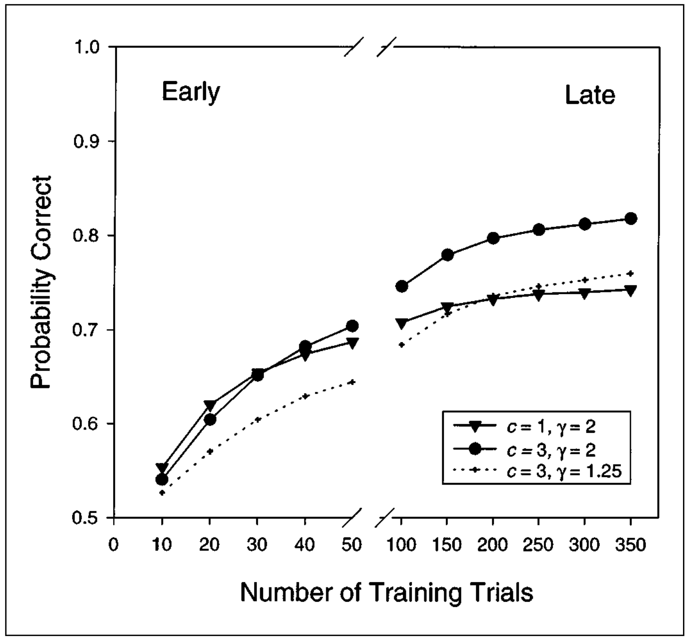

Fig. 3. Exemplar-model predictions of performance in the probabilistic classification learning task (Study 2) for low ($c = 1$) and high ($c = 3$) settings of the memory-sensitivity parameter (solid-line curves) and for a low setting ($\gamma = 1.25$) of the response-scaling parameter (dotted-line curve). Following the graphing procedure used by Knowlton et al. (1994, Fig. 2), we used an expanded scale along the horizontal axis for displaying the predictions for the first 50 trials of learning (early) and a compressed scale for the final 300 trials (late).

LIMITATIONS AND FUTURE DIRECTIONS

A Double Dissociation Involving Categorization and Recognition

We have addressed only a subset of the findings advanced by Knowlton, Squire, and their colleagues regarding the view that categorization and recognition are mediated by separate memory systems. In addition to testing control and amnesic subjects, for example, Knowlton et al. (1996) tested two groups of Parkinson's disease (PD) patients. Whereas the amnesics had shown impaired memory but intact early classification performance relative to normal participants, both groups of PD patients showed the reverse dissociation. The memory-questionnaire performance of the PD patients was as good as that of the normal participants, but their early classification learning was significantly worse than that of both the normal and the amnesic participants. (Late in classification learning, however, the PD patients eventually caught up to the amnesics.)

This double dissociation does indeed provide evidence in favor of the separate-memory-systems view; however, we believe that the issues involved are extremely complex. As indicated by our formal modeling, at least part of the double dissociation does not appear to be diagnostic. A single-system exemplar model that assumes that amnesics have lower memory sensitivity than do normal individuals predicts that amnesics will perform about as well as normal individuals during early classification. The main remaining question, therefore, is why the PD patients, who displayed good memory, performed worse than either group during early classification. This question has innumerable possible answers, many of which are unrelated to the issue of whether stored exemplars are used to represent categories. For example, note that a low setting of the response-scaling parameter γ in the exemplar model can be extremely detrimental to probabilistic classification. Optimal performance in the probabilistic classification task is achieved by using deterministic response rules (high values of γ), whereas suboptimal performance arises when γ is low. Observers with high memory sensitivity (large values of c) but suboptimal classification response strategies (low values of γ) could conceivably achieve excellent performance in a memory test, but exceptionally poor performance in probabilistic classification learning. The dotted learning curve in Figure 3 illustrates the exemplar-model predictions when there is a high value of memory sensitivity ($c = 3$) but the observer uses nearly a probability-matching response strategy ($\gamma = 1.25$). In this case, early classification learning is impaired, but the late-learning performance eventually catches up with that of the group that displays low memory sensitivity. This pattern is the one reported by Knowlton et al. (1996) regarding the relation between the PD patients and the amnesics.

We do not claim that PD patients use suboptimal, probabilistic response rules, and the precise basis for their poor classification performance is an important issue for future research.[5] Rather, we offer this example to illustrate that even a double dissociation does not provide clear-cut evidence of the involvement of separate memory systems. Multiple component processes interacting in nonlinear fashion in a single-memory-system model can produce such results as well.

5. Interestingly, Reber and Squire (1997) recently reported that a group of PD patients who performed poorly on the probabilistic classification task exhibited normal performance on a variety of other categorization tasks (such as prototype abstraction and artificial-grammar learning). These results support the view that the PD patients' deficit in probabilistic classification may reflect a strategic factor associated with that particular task (such as the type of response rule used) and not a generalized deficit in category learning and representation.

The Case of E.P.

Perhaps the greatest extant challenge for the single-system exemplar model is to account for data reported by Squire and Knowlton (1995) and Squire and Zola (1996) involving a single subject (E.P.) who suffers from profound anterograde and retrograde amnesia. Squire and Knowlton tested this subject in dot-pattern classification and recognition tasks similar to the ones we described and modeled earlier in this article. In a task in which a single dot pattern was presented 40 consecutive times during study, E.P.'s recognition performance was at chance, whereas control subjects performed with 95% accuracy. Impressively, however, E.P. performed with the same accuracy as did the control subjects in a categorization task in which 40 distinct dot patterns generated from a prototype had been presented during study. The baseline version of the exemplar model presented in this article is unable to fit this pattern of data.

Because the results pertain to a single subject and questions about generalizability arise, we are unsure how much weight to accord these data. Even E.P.'s behavior, however, may turn out not to be beyond the scope of an exemplar-based interpretation. For example, in the categorization task, roughly half of the 84 test items were members of the category and half were not. If E.P. had some residual exemplar-based category knowledge due to the original study episode, then the frequent presentations of new category members at the time of test could serve to redintegrate this knowledge. By contrast, in the test phase of the recognition task, there were only 8 presentations of the single study item and 76 presentations of distractor items. It is not inconceivable that, upon being flooded with distractors, an observer with extremely poor exemplar-based memory could suffer rapid interference or even "give up" and start guessing in the face of a seemingly impossible task.

Finally, we believe that E.P.'s data pose a challenge as well to Knowlton and Squire's hypothesis that categorization and recognition are mediated by separate memory systems. As reported by Squire and Knowlton (1995) and Squire and Zola (1996), in all tasks in which stimulus conditions were held constant, instructions to categorize versus recognize had relatively little effect on patterns of performance. For example, in the task in which 40 distinct patterns were presented during study, E.P. performed virtually the same at the time of test regardless of whether he was instructed to classify or recognize (Squire & Zola, 1996, Fig. 6). A similar result was observed for the control subjects. Such results seem consonant with the idea that categorization and recognition judgments are based on information from a single memory system.[6]

6. Indeed, to account for the results, Squire and Zola (1996) suggested that "when asked to recognize 40 stimuli that had been presented once each, both EP and control subjects tended to use a classification strategy" (p. 13518). Squire and Zola further suggested that "when asked to classify after seeing only one pattern 40 times, normal subjects tended to rely on declarative memory" (p. 13518). These suggestions begin to make murky the view that categorization is mediated by an implicit system and recognition by an explicit one.

SUMMARY

In summary, Knowlton, Squire, and their colleagues have reported a variety of intriguing dissociations between categorization and recognition performance. These dissociations are consistent with the view that separate memory systems underlie these cognitive tasks. In the present article, however, we demonstrated that various of the important dissociations are also apparently consistent with the idea that a single exemplar-based memory system underlies categorization and recognition, as long as one allows for plausible differences in parameter settings across groups. Whether or not a full account of Knowlton and Squire's complete set of dissociation findings is available within the exemplar approach remains a subject for future research, but we feel that we have made an important start.

Acknowledgments—This work was supported by Grant PHS R01 MH48494-07 and by Training Grant MH19879-05 (Modeling of Cognition) from the National Institute of Mental Health. The authors would like to thank Barbara Knowlton and Larry Squire for providing detailed information concerning the stimulus materials used in the studies of Knowlton and Squire (1993). The authors would also like to thank Asher Cohen, William Estes, Douglas Hintzman, Thomas Palmeri, David Shanks, Richard Shiffrin, Larry Squire, and three anonymous reviewers for helpful discussions or comments on earlier versions of this article.

REFERENCES

Anderson, J.R., Kline, P.J., & Beasley, C.M. (1979). A general learning theory and its applications to schema abstraction. In G.H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 13, pp. 277–318). San Diego: Academic Press.

D'Arcy, J.F., & Larochelle, S. (1997, May). *Dissociation of classification and recognition: A simulation with a single-store model.* Poster presented at the annual meeting of the American Psychological Society, Washington, DC.

Estes, W.K. (1986a). Array models for category learning. *Cognitive Psychology, 18,* 500–549.

Estes, W.K. (1986b). Memory storage and retrieval processes in category learning. *Journal of Experimental Psychology: General, 115,* 155–174.

Estes, W.K. (1994). *Classification and cognition.* New York: Oxford University Press.

Gillund, G., & Shiffrin, R.M. (1984). A retrieval model for both recognition and recall. *Psychological Review, 91,* 1–67.

Hintzman, D.L. (1986). "Schema abstraction" in a multiple-trace memory model. *Psychological Review, 93,* 411–428.

Hintzman, D.L. (1988). Judgments of frequency and recognition memory in a multiple-trace memory model. *Psychological Review, 95,* 528–551.

Knowlton, B.J., Mangels, J.A., & Squire, L.R. (1996). A neostriatal habit learning system in humans. *Science, 273,* 1399–1402.

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

Knowlton, B.J., Squire, L.R., & Gluck, M.A. (1994). Probabilistic classification learning in amnesia. *Learning and Memory, 1,* 106–120.

Maddox, W.T., & Ashby, F.G. (1993). Comparing decision bound and exemplar models of categorization. *Perception & Psychophysics, 53,* 49–70.

McKinley, S.C., & Nosofsky, R.M. (1995). Investigations of exemplar and decision bound models in large-size, ill-defined category structures. *Journal of Experimental Psychology: Human Perception and Performance, 21,* 128-148.

Medin, D.L. (1986). Comment on "Memory Storage and Retrieval Processes in Category Learning." *Journal of Experimental Psychology: General, 115,* 373–381.

Medin, D.L., & Schaffer, M.M. (1978). Context theory of classification learning. *Psychological Review, 85,* 207–238.

Metcalfe, J., & Fisher, R.P. (1986). The relation between recognition memory and classification learning. *Memory & Cognition, 14,* 164–173.

Nosofsky, R.M. (1984). Choice, similarity, and the context theory of classification. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 10,* 104–114.

Nosofsky, R.M. (1987). Attention and learning processes in the identification and categorization of integral stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13,* 87–108.

Nosofsky, R.M. (1988). Exemplar-based accounts of relations between classification, recognition, and typicality. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 14,* 700–708.

Nosofsky, R.M. (1991). Tests of an exemplar model for relating perceptual classification and recognition memory. *Journal of Experimental Psychology: Human Perception and Performance, 17,* 3–27.

Nosofsky, R.M. (1992). Exemplar-based approach to relating categorization, identification, and recognition. In F.G. Ashby (Ed.), *Multidimensional models of perception and cognition* (pp. 363–393). Hillsdale, NJ: Erlbaum.

Nosofsky, R.M., Kruschke, J.K., & McKinley, S.C. (1992). Combining exemplar-based category representations and connectionist learning rules. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 18,* 211–233.

Omohundro, J. (1981). Recognition vs. classification of ill-defined category exemplars. *Memory & Cognition, 9,* 324–331.

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

Reber, P.J., & Squire, L.R. (1997, November). *Implicit learning in patients with Parkinson's disease.* Poster presented at the annual meeting of the Psychonomic Society, Philadelphia.

Shanks, D.R., & St. John, M.F. (1994). Characteristics of dissociable human learning systems. *Behavioral and Brain Sciences, 17,* 367–447.

Shepard, R.N. (1987). Toward a universal law of generalization for psychological science. *Science, 237,* 1317–1323.

Shin, H.J., & Nosofsky, R.M. (1992). Similarity-scaling studies of dot-pattern classification and recognition. *Journal of Experimental Psychology: General, 121,* 278–304.

Squire, L.R., & Knowlton, B.J. (1995). Learning about categories in the absence of memory. *Proceedings of the National Academy of Sciences, USA, 92,* 12470–12474.

Squire, L.R., & Zola, S.M. (1996). Structure and function of declarative and non-declarative memory systems. *Proceedings of the National Academy of Sciences, USA, 93,* 13515–13522.

(RECEIVED 11/17/97; ACCEPTED 1/6/98)