Journey to the Center of the Category: The Dissociation in Amnesia between Categorization and Recognition

Article *in* Journal of Experimental Psychology Learning Memory and Cognition · August 2001

DOI: 10.1037/0278-7393.27.4.984 · Source: PubMed

John Paul Minda
The University of Western Ontario

Journal of Experimental Psychology:
Learning, Memory, and Cognition
2001, Vol. 27, No. 4, 984–1002

0278-7393/01/$5.00 DOI: 10.1037//0278-7393.27.4.984

Journey to the Center of the Category: The Dissociation in Amnesia Between Categorization and Recognition

J. David Smith and John Paul Minda
State University of New York at Buffalo

The authors' theoretical analysis of the dissociation in amnesia between categorization and recognition suggests these conclusions: (a) Comparing to-be-categorized items to a category center or prototype produces strong prototype advantages and steep typicality gradients, whereas comparing to-be-categorized items to the training exemplars that surround the prototype produces weak prototype advantages and flat typicality gradients; (b) participants often show the former pattern, suggesting their use of prototypes; (c) exemplar models account poorly for these categorization data, but prototype models account well for them; and (d) the recognition data suggest that controls use a single-comparison exemplar-memorization process more powerfully than amnesics. By pairing categorization based in prototypes with recognition based in exemplar memorization, the authors support and extend other recent accounts of cognitive performance that intermix prototypes and exemplars, and the authors reinforce traditional interpretations of the categorization-recognition dissociation in amnesia.

Some descriptions of categorization suggest that humans average their exemplar experiences to derive the category's center or prototype, compare new items to it, and accept the items as category members if similar enough. Other descriptions of categorization suggest that humans store their exemplar experiences and use these encapsulated episodes as comparative standards in categorization. Either claim—that humans abstract generalities or store particulars—is an important claim about cognition. But the issue remains unresolved.

One reason for this is that both the prototype and exemplar theories are powerful enough to explain many phenomena. For example, early studies showed that the categorization of training exemplars decays more rapidly than the categorization of transfer items (especially prototypical transfer items; Homa, Cross, Cornell, Goldman, & Schwartz, 1973; Posner & Keele, 1970; Strange, Keeney, Kessel, & Jenkins, 1970). This result was first interpreted as showing that humans learn prototypes during category training and use them in transfer even while partially forgetting the training exemplars. But subsequent formal analyses suggested that a unitary, exemplar-based system could also explain this result (Hintzman, 1986; Hintzman & Ludlam, 1980).

Another example concerns the independence of categorization and recognition performance. Participants sometimes show better categorization than recognition for specific training exemplars (Metcalfe & Fisher, 1986; Omohundro, 1981) or low correlations between categorization and recognition performance (Hayes-Roth & Hayes-Roth, 1977). These results challenge the idea that both performances use the same exemplar traces and processes, because exemplars stored well enough to serve categorization should serve recognition also (Anderson, Kline, & Beasley, 1979; Metcalfe & Fisher, 1986; Omohundro, 1981). Of course, separate prototype-based and exemplar-based processes, serving categorization and recognition, respectively, explain this result. But subsequent formal analyses suggested that a unitary exemplar-based system could also explain this result (Nosofsky, 1988, 1991; Shin & Nosofsky, 1992).

Similarly, Nosofsky and Zaki (1998) simulated the dissociation between categorization and recognition performance in amnesia by studying both performances in control participants across a week's delay. Categorization performance dropped slightly (4%); recognition performance dropped sharply (17%). Nosofsky and Zaki explained these results by assuming a unitary system that compares the to-be-categorized or to-be-recognized item to the relevant stored exemplar traces. But they acknowledged that these results could also be explained by assuming multiple memory systems containing robust prototypes and forgettable exemplars.

Finally, there is the actual categorization–recognition dissociation shown by amnesics. Research by Knowlton and Squire (Knowlton, Mangels, & Squire, 1996; Knowlton & Squire, 1993; Knowlton, Squire, & Gluck, 1994) showed that amnesics perform relatively normally when categorizing dot patterns derived from an underlying prototype. But they are impaired in performing an old–new recognition task with similar materials. Indeed, Squire and Knowlton (1995) showed that one amnesic categorized relatively normally compared with age-matched controls even while showing no recognition at all. The amnesia dissociation challenges an exemplar theory that assumes a unitary processing system based in similarity comparisons to stored exemplars because the same informational pool serves one task but not the other. The amnesia dissociation is also problematic for exemplar theory because the formal account provided by Nosofsky of categorization–recognition independence (Nosofsky, 1988) does not apply to

J. David Smith, Department of Psychology and Center for Cognitive Science, State University of New York at Buffalo; John Paul Minda, Department of Psychology, State University of New York at Buffalo.

John Paul Minda is now at the Beckman Institute, University of Illinois at Urbana–Champaign.

Correspondence concerning this article should be addressed to J. David Smith, Department of Psychology, Park Hall, State University of New York, Buffalo, New York 14260. Electronic mail may be sent to psysmith@acsu.buffalo.edu.

Knowlton and Squire's (1993) experimental paradigm (Nosofsky & Zaki, 1998, p. 249). The amnesia dissociation is the focus of this article, though it bears more broadly on the roles that prototype and exemplar theories have in explaining cognitive performance.

Knowlton and Squire (1993) used multiple memory systems to explain their results. They assumed that categorization performance relies on an implicit memory system—intact in amnesics—that represents category-level information in the form of prototypes. They assumed that recognition performance relies on an explicit memory system—impaired in amnesics—that contains declarative memories about specific exemplars.

Nosofsky and Zaki (1998) took an alternative approach to this problem. They assumed that both controls and amnesics successfully store exemplars in memory. They argued that a unitary system based on similarities to those stored exemplars could predict the dissociation by assuming that amnesics use specific exemplar information less easily or sensitively than controls, leaving amnesics able to categorize using stored exemplars but unable to recognize using them. They captured these different levels of sensitivity with a parameter change between groups that changed the impact that exemplar comparisons have on the cognitive processes of controls and amnesics. In fact, Nosofsky and Zaki specified a system of formal equations that instantiated their unitary, exemplar-based processing system. They fit this model to the categorization and recognition performance of controls and amnesics.

Nosofsky and Zaki's (1998) exemplar-based interpretation of the categorization–recognition dissociation in amnesia is potentially important for several reasons. First, it suggests that even a strong categorization–recognition dissociation can be explained using only exemplar-based processes. Second, it raises questions about the appropriate interpretation of one of neuropsychology's influential results. Third, it challenges the idea of multiple memory systems. Fourth, it undermines an important result in the literature showing prototype abstraction. Fifth, the exemplar-based interpretation of Nosofsky and Zaki is instantiated so carefully that it can be analyzed in detail and possibly even disconfirmed. Accordingly, the present article shows that the exemplar-based interpretation of the categorization–recognition dissociation in amnesia can be disconfirmed. In fact, this dissociation provides strong evidence for prototype-based processes and multiple memory processes in human cognition.

The present article proceeds as follows. First, we describe the dot-distortion methodology used by Posner, Homa, Knowlton and Squire, Nosofsky and Zaki, and others, and we summarize the psychophysical principle that governs similarity among dot patterns. Second, we review Knowlton and Squire's tasks and results that showed the amnesia dissociation. Third, we describe the distinctive power model that Nosofsky and Zaki used to model the categorization performance of controls and amnesics. Fourth, we show that this model's distinctive features make it theoretically ambiguous and even similar to prototype models. Fifth, therefore, we describe an alternative approach that lets the dot-distortion paradigm clearly disambiguate prototype-based and exemplar-based processes. Sixth, formalizing this approach, we show that a variety of exemplar models (including more traditional ones), when fit to Knowlton and Squire's dot-distortion data, qualitatively fail to describe what controls and amnesics do in the categorization task. In contrast, an equivalent prototype model fits the data well. Seventh, we address the recognition performance of Knowlton and Squire's participants. Here we consider whether participants (especially controls) simply memorize the exemplars in the study phase of the recognition task, or whether they use the systematic exemplar-to-exemplar comparisons assumed by exemplar-based approaches. We find strong support for exemplar memorization in the data but no support for the exemplar-comparison process assumed by exemplar theory.

Background

Stimulus Materials

The stimulus materials of Knowlton and Squire (1993) were created with a well-established method that generates families of dot patterns from prototypes. In this method, nine points are randomly selected from within the central 30×30 area of a 50×50 grid. These nine dots are a prototype. The distortions (the family members) are produced by applying a series of probabilities that determine whether each dot will keep the same position as it did in the prototype, move to the first layer of 8 grid squares around the original position, move to the second layer of 16 grid squares, or move farther away (Posner, Goldsmith, & Welton, 1967, Table 1, p. 31). By using different series of probabilities, one can arrange for dot patterns that are low-, medium-, or high-level distortions of the prototype and that form families of dot patterns that have strong, moderate, or weak family resemblance to the prototype. These different families of distortions can be thought of as lying on hyperspheres (N-dimensional shells) of different radii surrounding the central prototype, with the radius determined by the distortion level (Homa, Dunbar, & Nohre, 1991; Homa, Sterling, & Trepel, 1981).

The Psychological Similarity Between Dot Patterns

Posner and his colleagues (Posner, 1964; Posner et al., 1967) carried out a study of the similarity among dot patterns as perceived and rated by hundreds of human observers. We summarize this research because it has important applications in the present article and implications for evaluating exemplar theory.

In one study, Posner et al. (1967) focused on prototype-based similarity and distance. They asked participants to rate the dissimilarity between prototypes and their distortions. Subjective dissimilarity increased linearly as prototypes and higher level distortions were compared (see Posner et al., 1967, Table 2, p. 32, and Figure 3, p. 31; see Figure 1A in the present article). Posner et al. were also able to calculate the objective physical distance that the nine dots in a pattern were moved on average when a distortion was created from its prototype. They found that logarithmic objective distance also increased linearly with distortion level (Figure 1B in the present article). Combining these results, Posner et al. concluded that participants' subjective similarity ratings were a linear function of the logarithmic objective distance that corresponding dots moved between patterns. This relationship was so strong (a correlation of .99 across the four similarity–distance pairs in Figures 1A and 1B) that all the variance in subjective similarity was accounted for by objective distance.

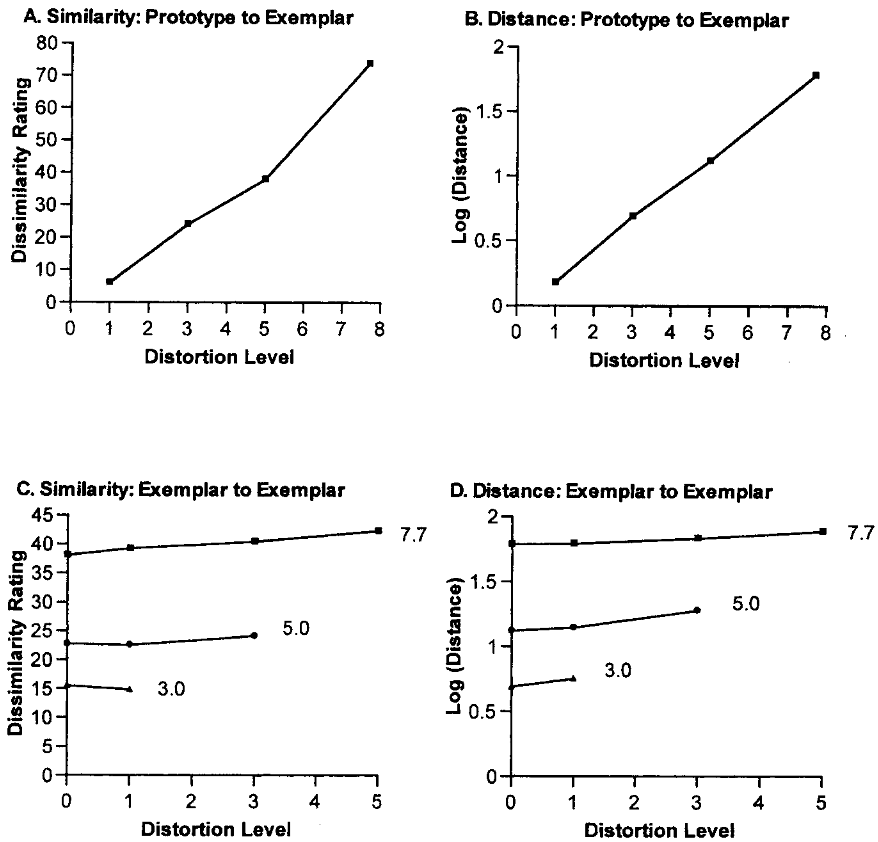

Figure 1. Panel A: The median dissimilarity rating provided by observers in Posner et al.'s (1967) study, for patterns at different levels of distortion compared to their prototype. Distortion level was defined as the amount of uncertainty (bits per dot) regarding the positioning of the dots in a pattern. The similarity-rating data were taken from Posner et al. (1967, Table 2, p. 32) by averaging across the five prototypes used in the experiment. Panel B: The relationship between the measure ln (1 + average Pythagorean distance moved per dot) and distortion level (as defined by Posner et al. and many others including Knowlton & Squire, 1993, and Nosofsky & Zaki, 1998). This relationship is illustrated with a simulation that created thousands of prototype–distortion pairs (with equal numbers of distortions at Level 1, Level 3, Level 5, or Level 7.7) and found the average logarithmic interdot distance for each distortion level. Panel C: The median dissimilarity rating provided by observers in Posner et al.'s study when patterns at three levels of distortion (Level 7.7, Level 5, and Level 3) were compared to all lower levels of distortion (e.g., Level 7.7 distortions compared to Level 5, Level 3, Level 1, and Level 0 distortions). Panel D: The mean of the measure ln (1 + average Pythagorean distance moved per dot) when patterns at three levels of distortion (Level 7.7, Level 5, and Level 3) were compared to all lower levels of distortion (e.g., Level 7.7 distortions compared to Level 5, Level 3, Level 1, and Level 0 distortions). These means are based on a simulation that created thousands of distortion–distortion pairs and found the average logarithmic interdot distance within each combination of distortion levels (e.g., Level 7.7–Level 5, Level 3–Level 1).

In another study, Posner et al. (1967) focused on exemplar-based similarity and distance. They asked participants to compare distortions at a given level with distortions at lower levels. In this case they found flat similarity-scaling functions (see Posner et al., 1967, Table 3, p. 34). This result is illustrated in Figure 1C. It shows the median subjective dissimilarity ratings provided by Posner et al.'s observers for three distortion levels (7.7 bits/dot, 5 bits/dot, and 3 bits/dot) compared to all lower distortion levels (e.g., 7.7 compared to 5; 7.7 to 3; 7.7 to 1; and 7.7 to 0). Subjective dissimilarity barely decreases at a given distortion level as the comparison item moves in toward the prototype of the category (i.e., as the comparison item becomes a progressively lower level distortion). This result was explained when Posner et al. also found flat objective-distance functions when they compared distortions at any level to distortions at lower levels (see Posner et al., 1967, Table 4, p. 34). That is, objective interdot distance also barely decreased as the inner stimulus pattern moved in toward the category's center. This result is illustrated in Figure 1D, again using a logarithmic distance measure.

These results confirmed that subjective similarity was a linear function of objective distance. Again this relationship was so strong (a correlation of .99 across the nine similarity–distance pairs in Figures 1C and 1D) that distance captured all the variance in similarity. But here both distance and similarity were controlled by the level of the outer (higher level) distortion—that is, by the radius of the more distant exemplar shell. Neither distance nor similarity changed much as comparison items moved inside that exemplar shell toward the category's center.

Re-Creating Posner et al.'s (1967) Result

The flatness of these exemplar-based similarity gradients has important implications for evaluating exemplar theory. Accordingly, we conducted a new experiment that confirmed Posner et al.'s (1967) results and broadened the basis for their claim. To do so, we repeated their experiment using the random-dot polygons that have been the other important class of stimulus materials arising from the dot-distortion methodology. (Random-dot polygons are simply random-dot patterns with the dots joined to form polygons.) In our experiment, 37 participants rated 480 pairs of random-polygon patterns for their dissimilarity on a scale of 1 (*no difference*) to 6 (*big difference*). The two stimuli on each trial were always 0-, 1-, 3-, 5-, or 7.7-bits-per-dot distortions of the same prototype, and each participant rated 32 instances of every pairwise combination of these (0-0, 0-1, . . . , 7.7-7.7). The prototype was chosen randomly on every one of the 37 × 480 trials, making this experiment a comprehensive survey of dot-distortion space.

Figure 2 shows our results. Subjective dissimilarity and objective distance between a prototype and its distortions increased linearly with distortion level (Figures 2A and 2B). The correlation between dissimilarity and distance was .999 across these four dissimilarity–distance pairs. Figure 2 also shows the flatness of the similarity-scaling functions when a distortion at a given level was compared to distortions that were closer to the prototype. Subjective dissimilarity (Figure 2C) and objective distance (Figure 2D) barely decreased. The correlation between dissimilarity and distance over these nine dissimilarity–distance pairs was .998. Both correlations (.999 and .998) confirm that the variation in subjective similarity was completely explained by the variation in objective distance.

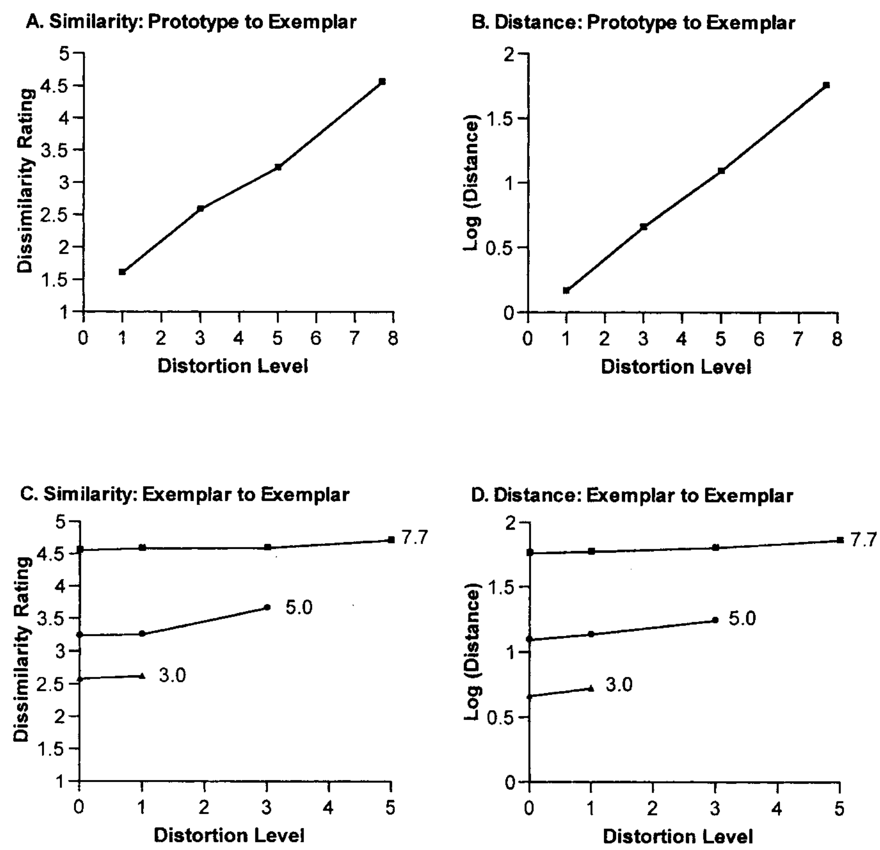

Figure 2. Panel A: The mean dissimilarity rating provided by observers for patterns at different levels of distortion from their originating prototype. These data are from the present experiment that re-created the results of Posner et al. (1967). Panel B: The relationship between the measure ln (1 + average Pythagorean distance moved per dot) and distortion level for the stimuli used in the replicating experiment. Panel C: The mean dissimilarity rating provided by observers in the replicating experiment when patterns at three levels of distortion (Level 7.7, Level 5, and Level 3) were compared to all lower levels of distortion (e.g., Level 7.7 distortions compared to Level 5, Level 3, Level 1, and Level 0 distortions). Panel D: The mean logarithmic distance between patterns in the replication experiment when patterns at three levels of distortion (Level 7.7, Level 5, and Level 3) were compared to all lower levels of distortion (e.g., Level 7.7 distortions compared to Level 5, Level 3, Level 1, and Level 0 distortions).

These results endorse the psychophysics of dot distortions that Posner et al. (1967) established. This psychophysics is a lasting contribution that has an important application in the present article. Posner et al. showed—across samples, across patterns, across distortion algorithms (Posner, 1964), and for both prototype-centered distances (Figures 1A and 1B) and exemplar-centered distances (Figures 1C and 1D)—that the subjective similarity ratings that observers provide in pairwise comparisons of dot patterns have a precisely linear relation to the logarithm of the average Pythagorean distance moved per dot.

The simplicity of this relation allows easy and accurate interpolation and extrapolation to all required points along the function relating similarity to distance. The relation naturally provides measures of psychological similarity that can become the input to formal models of categorization. Moreover, Posner et al. (1967) established this relationship in a neutral task context regarding learning or instructional sets. They established it in a neutral theoretical context regarding prototypes and exemplars in categorization, because these studies predated this debate in the literature. Indeed, Posner et al. found that exactly the same similarity–distance relationship was applicable to both prototype-based similarities and exemplar-based similarities. In the present article we incorporate this logarithmic distance measure that summarizes psychological similarity so simply and evenhandedly.

Tasks

Categorization Task

Knowlton and Squire's (1993) participants viewed 40 high-level distortions of a prototype for 5 s each. They were told that all of the patterns belonged to a single category. Five minutes later, participants judged whether 84 transfer dot patterns were members of the training category. The 84 transfer patterns comprised 4 presentations of the prototype that was never seen in training, 20 new low-level distortions of the prototype, 20 new high-level distortions, and 40 or random unrelated patterns outside the category.

Recognition Task

Knowlton and Squire's (1993) participants viewed the five members of the recognition list eight times for 5 s each. The list members were each a high-level distortion of a randomly selected prototype. Five minutes later, they judged whether each of 10 test items was exactly the same as one from the recognition list. The 10 test patterns comprised the 5 list patterns and 5 new patterns that were each a high-level distortion of a randomly selected prototype.

Results

The crucial data are 10 performance levels—the levels of category endorsement by controls and amnesics for prototypes, low-level distortions, high-level distortions, and random or unrelated patterns in the category task, and the levels of accuracy that controls and amnesics achieved in the recognition task. These data are shown for both groups in Figure 3. Controls and amnesics categorized about the same for each of the four item types. But controls performed more accurately in the recognition task because they distinguished old from new items more sensitively than amnesics. This is the categorization–recognition dissociation in amnesia.

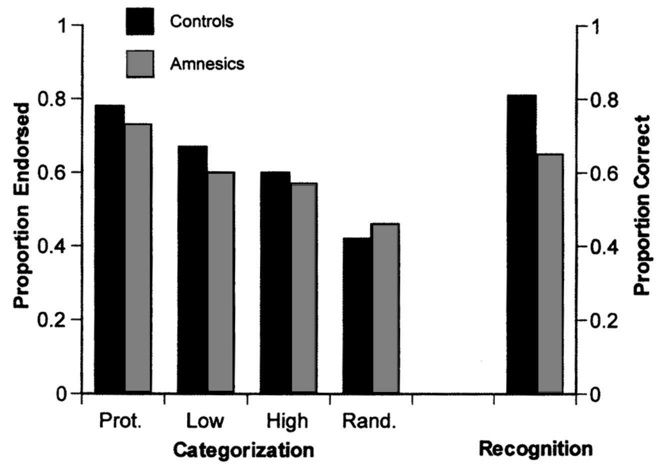

Figure 3. The categorization and recognition performance of control and amnesic participants from Knowlton and Squire (1993). Prot. = prototype; Rand. = random.

Theoretical Analyses

Categorization

The Power Model of Categorization

Nosofsky and Zaki (1998) instantiated an exemplar-based power model that they thought could explain the amnesia dissociation. Like exemplar models, the power model tries to estimate the total similarity a to-be-categorized item type (prototype, low-level distortion, high-level distortion, or random) would have to the 40 high-level distortion training exemplars. It then uses this total similarity to estimate the strength of category endorsement that would result for that item type. However, Nosofsky and Zaki's power model is distinctively different from traditional exemplar models in ways that are theoretically important. Thus, we describe it carefully here.

In the model, the probability of endorsing an item (i) into Training Category A—that is, the probability of Response A (R_A) given Stimulus i (S_i)—is

$$P(R_A|S_i) = \frac{40 \times sr(i, h)^p}{40 \times sr(i, h)^p + k_{40}}. \quad (1)$$

This equation estimates (in its numerator) the total similarity an item type has to the 40 high-level distortions used in training. It then compares this total similarity with a criterion quantity (k_{40}) to judge the extent of category endorsement by the extent to which the criterion threshold is exceeded.

The sr quantities in this equation come from a similarity-rating task that Nosofsky and Zaki (1998) conducted with their participants after they had completed both category training and testing. Participants rated the subjective similarity between training distortions and the prototype, low-level distortions, high-level distortions, and random items used in testing. The quantity $sr(i, h)$ thus denotes the average similarity rating (sr) participants gave between some item type (i) presented in testing and the high-level distortions (h) presented in training. This average was taken over participants and over test item–training item pairs.

Nosofsky and Zaki (1998) further assumed that these similarity ratings did not directly represent psychological similarity. Instead, they assumed that the psychological similarity between types of items was given by a power transformation of their rated similarity. In accordance, Nosofsky and Zaki raised these averaged similarity ratings to a power (p) that was a free parameter in the power model. For example, given an average similarity rating of 5.311 between training distortions and prototype test items, Nosofsky and Zaki estimated that the psychological similarity between these item types was $5.311^{5.18}$ (5,707) for controls. In contrast, they estimated that the psychological similarity between these item types was $5.311^{2.99}$ (147) for amnesics. This power transformation let Nosofsky and Zaki grant control and amnesic participants differential sensitivity in the category task. It was an appropriate way for them to instantiate their basic hypothesis that amnesics store exemplars as controls do but then access those exemplars less sensitively. Nosofsky and Zaki made the further decision to multiply the averaged similarity ratings by 40. This scaled these averages up to the level expected as a total, had an item been actually compared to the 40 high-level distortions from training.

Finally, Nosofsky and Zaki (1998) used the free parameter k_{40} as a criterion quantity that provides a proportionalizing threshold to the model. (Thus, the power model of categorization has two free parameters.) That is, the estimated total similarity of an item type to 40 of the training exemplars (the choice rule's numerator) must exceed this threshold before the item type can be endorsed as belonging in the category beyond 50%. The more the total category similarity exceeds the threshold, the more strongly the item type will be endorsed. The use of a criterion quantity was another departure from traditional exemplar models. However, it was a natural one because Knowlton and Squire's (1993) experiment had no contrast category so there was no total contrast-category similarity to complete the denominator of the choice rule in the usual way. The quantity k_{40} can constructively be interpreted as the total of 40 null or noncategory similarity impacts on the cognitive system. The power model's decision rule thus examines how strongly 40 average category similarities exceed 40 average noncategory similarities.

An Averaged Model

Exemplar models typically use specific item-to-item similarities in predicting the categorization performance for every individual item in the task. Regarding the present category task, these models would compare each of the 84 transfer items to all of the 40 stored training exemplars and predict 84 levels of categorization performance. This fine-grained analytic approach would be true to exemplar theory's assumptions of specific exemplar storage and systematic similarity comparisons using these exemplar representations.

In contrast, the power model of categorization is distinctive for being an averaged model. The training items are treated as average high-level distortions; the transfer items are treated as an average token of some stimulus type. Averaged similarity ratings are input into the model to predict the average categorization level for each of the four stimulus types.

There were valid reasons for proceeding in this way. The data from dot-pattern studies are usually reported and interpreted as averages. The Knowlton and Squire (1993) data were reported and interpreted as averages. Even though Nosofsky and Zaki (1998) did collect the relevant pairwise ratings that could have supported more fine-grained analyses, the number of stimulus pairs (40 × 84), spread over limited participants limited in the number of pairs they can rate, might have left the individual item-to-item similarity estimates too weak to support more detailed predictions. However, we now show that the averaged character of the model creates a theoretical ambiguity about categorization performance that leads one to reconsider exemplar-based interpretations.

Forty Comparisons or One

The power model's categorization equation,

$$P(R_A|S_i) = \frac{40 \times sr(i, h)^p}{40 \times sr(i, h)^p + k_{40}}, \tag{2}$$

can be rewritten as

$$P(R_A|S_i) = \frac{40 \times sr(i, h)^p}{40 \times sr(i, h)^p + 40 \times k_{40}/40} \tag{3}$$

or as

$$P(R_A|S_i) = \frac{40 \times sr(i, h)^p}{40 \times sr(i, h)^p + 40 \times k_1} \tag{4}$$

and simplified to

$$P(R_A|S_i) = \frac{sr(i, h)^p}{sr(i, h)^p + k_1}. \tag{5}$$

The simplified equation shows that it need not be that 40 training items were stored and 40 similarities summed to make a categorization decision. We could also replace the 40s with 10s, but this would not imply that participants store every fourth exemplar and sum 10 similarities in order to categorize an item. Neither the 40s nor 10s nor their absence changes the formal equation. Neither the 40s nor 10s nor their absence makes anything true or not true about minds. But when one eliminates the extras, the appearance of systematic comparisons to all the training exemplars is eliminated. The equation is just a ratio of some trained category similarity ($sr(i, h)^p$) to some criterion or null similarity (k_1). It is only about the resonance between the to-be-categorized item and some category information compared with some baseline resonance. The equation leaves open the theoretical possibilities that the separateness of the 40 items has been lost in memory, that they have melded or merged into a single, general representation, that category decisions refer the to-be-categorized item to this single representation, that this single representation is a prototype. One does not have to prefer this interpretation of the power model's categorization equation, only see that it is viable.

Shared Characteristics of Averaged Models and Prototype Models

Other aspects of the power model of categorization make plain its theoretical ambiguity. It has no specific-exemplar parameter (unlike the recognition model that is its counterpart)—this parameter seemed unnecessary because Knowlton and Squire (1993) only transferred their participants to new items. Unlike other exemplar models, therefore, it provides no way for old exemplars to retrieve themselves in memory and receive a boost in categorization accuracy. Thus, it cannot predict the old-exemplar advantage that is one hallmark of the success of exemplar models (Smith & Minda, 2000). It cannot account for the exemplar-density effects that are another hallmark of exemplar theory. It has no way to master categories that are not linearly separable (e.g., categories that have exception items or that instantiate the XOR problem). In all these ways the averaged power model has closer links to traditional prototype models than to traditional exemplar models with their individuated exemplar representations and their systematic exemplar-to-exemplar comparisons.

The strong conclusion from this would be that the power model of categorization shows that controls and amnesics categorize to about the same level using processes grounded more in the category's general statistical tendencies than in the individuated exemplars and systematic comparisons that usually underlie exemplar models. This interpretation would echo the traditional interpretation that controls and amnesics are both able to use category-general information in the service of categorization. The mild conclusion from this would be that categorization by controls and amnesics remains ambiguous regarding its representational basis in prototypes or exemplars.

However, in this case the situation still needs clarifying. To do so, one needs a way to distinguish averaged models from prototype models. One needs a way to distinguish exemplar-based categorization performances from prototype-based categorization performances.

Distinguishing Exemplar-Based and Prototype-Based Similarity

Posner et al.'s (1967) research suggests a way to make this distinction (see Figures 1 and 2 in the present article). Examining our figures, one notes an important asymmetry. As dot patterns are made that are less distorted, they decrease strongly in dissimilarity to the prototype (Figures 1A and 2A) but they barely decrease in dissimilarity to a shell of exemplars (Figures 1C and 2C). Because this asymmetry provides a critical insight about the geometrical properties of exemplar and prototype theories, and about exemplar and prototype models, we explain it intuitively, illustrate it in a two-dimensional space that preserves geometrical intuitions, and illustrate it in the nine-dimensional dot-distortion space of interest here.

As an analogy, consider a distant spaceship moving toward a sun (prototype) with a shell of planets (exemplars) around it, and imagine plotting at each moment the distance of the spaceship from the sun and its average distance from the planets. At first the spaceship will approach the whole solar system simultaneously, and the distances to the sun and to the planets will decrease together. However, as the spaceship nears the sun and moves inside the shell of planets, the asymmetry observed by Posner et al. (1967) must occur. Now the spaceship's distance from the sun will still decrease. But now the spaceship will only get closer to a few planets on the opposite side of the system. It will move past other planets off to each side without changing distance to them. It will even move away from planets on the side of the system it entered. Thus, planet-centered distance will approach a steady state of balanced decreasing, unchanging, and increasing distance, even while sun-centered distance still decreases.

To establish this relationship systematically in a low-dimensional space, we defined a prototype as a single XY-coordinate pair at the center of a circle (to give it a central value on two continuous dimensions). Then we randomly chose 40 training exemplars as single XY-coordinate pairs on a circular shell surrounding it, so that each of them represented the same-level distortion of the prototype's central values. Then we chose a single stimulus and found its distance from the prototype center and its average distance from the 40 training exemplars on the shell. We repeated the process of choosing the 40 training exemplars and the stimulus 1,920 times, choosing stimuli that systematically moved in toward the prototype center and the exemplar shell.

Figure 4A shows the expected disconnect between prototype-

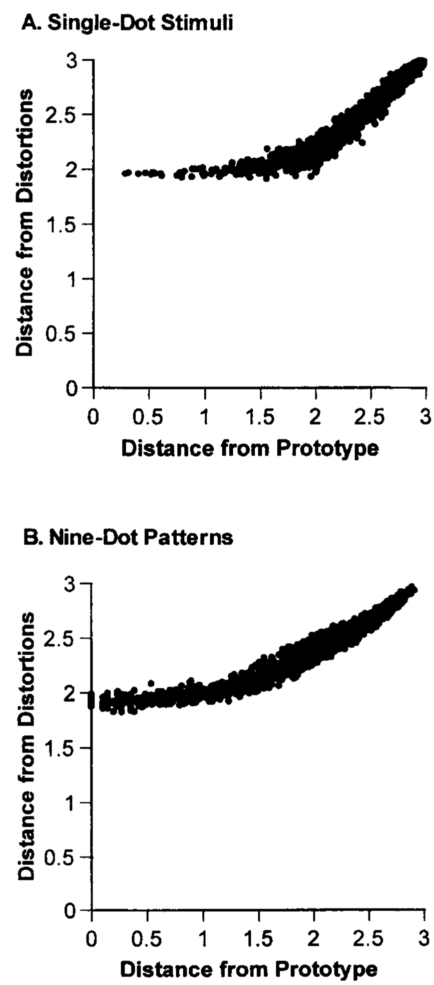

Figure 4. Panel A: Relation between distance from a circle's center and distance from 40 randomly chosen points on the circle for 1,920 points that systematically moved in toward the circle's center. Panel B: Relation between distance from a prototype and distance from 40 randomly chosen exemplars in a shell around it, for 1,920 random-dot patterns that systematically moved in toward the prototype.

based and exemplar-based distance that is strongest nearest the prototype. Starting from the right and moving leftward along the x-axis, at first the two distances decrease together as the far-away stimulus moves closer to the prototype and the exemplar shell. Near-in stimuli continue to decrease in distance to the prototype but do not decrease in distance to the exemplar shell.

To establish this relationship systematically for dot patterns, we chose a random prototype to represent a category center and we randomly chose 40 distortions lying a known distance away from the prototype to represent its training exemplars. These 40 training exemplars lie on a shell surrounding the prototype; the shell's radius is the known distance given by the distortion level that is chosen for the shell (7.7 bits/dot for this simulation). Then we chose a single stimulus and found its distance from the prototype and its average distance from the 40 training exemplars. This entire process was repeated 1,920 times, choosing single stimuli that systematically moved in toward the prototype center and the exemplar shell. Figure 4B shows the expected result. Prototype-based distance decreases strongly all the way to the center of the category. In contrast, exemplar-based distance reaches the steady state that Posner et al. (1967) discussed, that their observers showed in their similarity ratings (Figure 1C), and that our observers showed as well (Figure 2C).

To illustrate the generality of the difference between exemplar-based similarity and prototype-based similarity, we calculated in different ways the distance between Knowlton and Squire's (1993) 40 training exemplars and their 84 transfer items. In different cases, we assumed that participants focused their attention on one dot, on a three-dot corner, or on a five-dot constellation. In other cases, we assumed that the similarity space was organized by a city-block metric (Minkowski Metric 1), by a Euclidean metric (Minkowski Metric 2), or by Minkowski Metric 3. Figures 5A and 5B show that the geometries of exemplar-based comparisons and prototype-based comparisons transcend different attentional strategies and different similarity metrics. Exemplar-based comparisons always create flat typicality gradients close in to the center of the category and small prototype effects. Prototype-based comparisons always create the opposite pattern.

It is important to realize that this difference has a geometrical and mathematical necessity that will cause it to hold over almost all possible psychologies of dot patterns. The 40 training exemplars have been strongly, randomly distorted away from the prototype. This ensures that—under any psychological description of the stimuli—they will be spread out over a region of multidimensional similarity space. This spread will make it impossible for a stimulus to be close to all 40 training exemplars simultaneously. It will ensure that all the stimuli inside the space of the exemplars will be like some patterns and unlike others. If the exemplars really are the comparative and reference standard for the category, all the stimuli inside the shell of the exemplars (including low-level distortions and the prototype) will be about equally similar to the training exemplars and so the similarity gradients in the task will be quite flat. In contrast, stimuli will be able to close in on the prototype because it is a single point in the space, not 40 scattered points. If the prototype really is the comparative and reference standard for the category, this will allow some very high-similarity stimuli and it will produce steep typicality gradients.

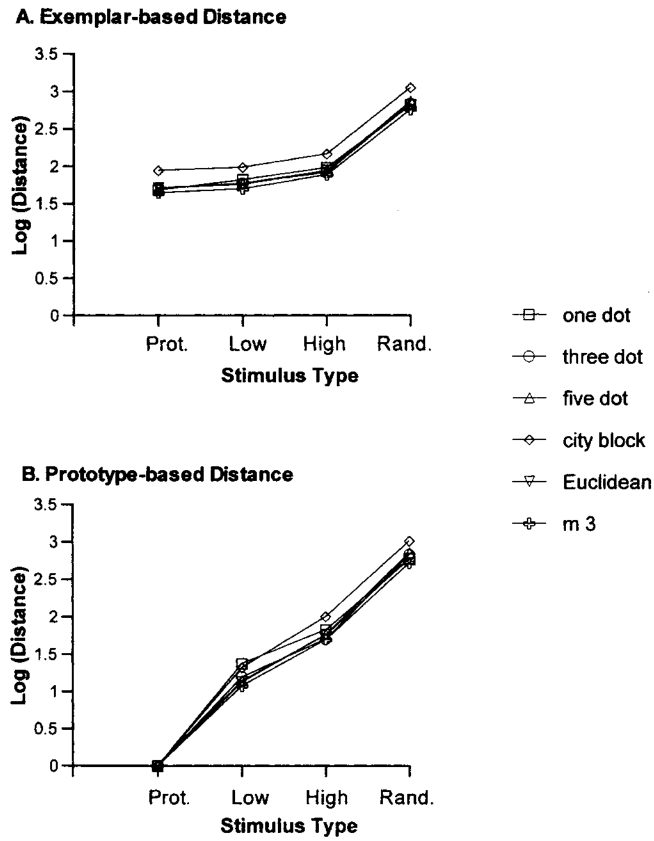

Figure 5. Panel A: The average logarithmic distance between Knowlton and Squire's (1993) prototype, low-level distortion, high-level distortion, and random test stimuli and their 40 high-level distortion training items, calculated assuming that participants focused their attention on one dot, on a three-dot corner, or on a five-dot constellation, and also calculated assuming that similarity space was organized according to a city-block metric (Minkowski Metric 1), by a Euclidean metric (Minkowski Metric 2), or by Minkowski Metric 3. Panel B: The average logarithmic distance between Knowlton and Squire's prototype, low-level distortion, high-level distortion, and random test stimuli and their prototype, calculated assuming that participants attended to one, three, or five dots, and calculated assuming each of three Minkowski metrics. Prot. = prototype; Rand. = random.

Distinguishing Exemplar-Based and Prototype-Based Categorization Performance

This difference has implications for theories and models of categorization and for interpretations of observed categorization performance. These implications hold for the range of theories and models that translate distances (exemplar based or prototype based) into psychological similarities and these into predicted response levels.

First, to-be-categorized items will have a narrower range of distances from the family of training exemplars than from the prototype. The range restriction applies in turn to psychological similarities and to predicted response levels. Consequently, exemplar-based models will tend to predict homogeneous performance profiles and flat typicality gradients. If one observes such a profile, a possible reason for it will be that participants were using exemplar-based processing in making categorization decisions. In contrast, prototype models will tend to predict heterogeneous performance profiles and sharp typicality gradients. If one observes such a profile, a possible reason for it will be that participants were using prototype-based processing in making categorization decisions.

Second, this range restriction is most dramatic for to-be-categorized items that are typical category members and close to the prototype. If the cognitive system does compare to-be-categorized items to the shell of training exemplars that surround the prototype, then prototypes, highly typical members, and typical members will all seem about equally distant from the training shell, equally similar to the training exemplars, and equally strongly in the category. In particular, exemplar models will predict small performance advantages for prototypes over other stimulus types. In contrast, prototype models will predict strong prototype-item advantages because they operate under no range restriction.

These facts can be illustrated in the same two domains as before. First, we defined two single-dot prototypes in two-dimensional Euclidean space (100, 100; 150, 150) and built categories containing the prototype of each category (Stimuli 1 and 9), six items at a radius of 10 from each prototype (Stimuli 2–7, 10–15), and one item at a radius of 40 from each prototype (Stimuli 8 and 16). Then we used standard multiplicative prototype and exemplar models to find what general predictions the two classes of model make about performance levels for each of the 16 stimuli. The choice rules for these prototype and exemplar models, respectively, were as follows:

$$P(R_A|S_i) = \frac{\eta_{iP_A}}{\eta_{iP_A} + \eta_{iP_B}} \tag{6}$$

or

$$P(R_A|S_i) = \frac{\sum\limits_{j\epsilon C_A} \eta_{ij}}{\sum\limits_{j\epsilon C_A} \eta_{ij} + \sum\limits_{j\epsilon C_B} \eta_{ij}}. \tag{7}$$

For the prototype model, the level of Category A responses predicted for a to-be-categorized item (i) was taken to be its similarity (η) to Prototype A (P_A) divided by its combined similarity to both prototypes. For the exemplar model, the level of Category A responses predicted was taken to be the summed similarity to the eight Category A members divided by the combined similarity to both exemplar sets. These equations are closely analogous to those for the power model except that (as is standard) the criterion quantity in the denominator has been replaced by the total similarity the item had to the prototype or the exemplars of the second, contrast category.

The similarity between any two items in both models was calculated as follows. We began with the logarithmic interdot distance between the items. As is standard, distance was then scaled by a sensitivity level (the only free parameter in these two models). As is standard, this scaled distance was transformed into a similarity with an exponential-decay function.

For each model, we chose 5,000 random sensitivities, found the 16 predicted performance levels for that configuration of the model, and averaged the 5,000 predicted profiles into the graphs shown in Figures 6A and 6B. The exemplar model produces

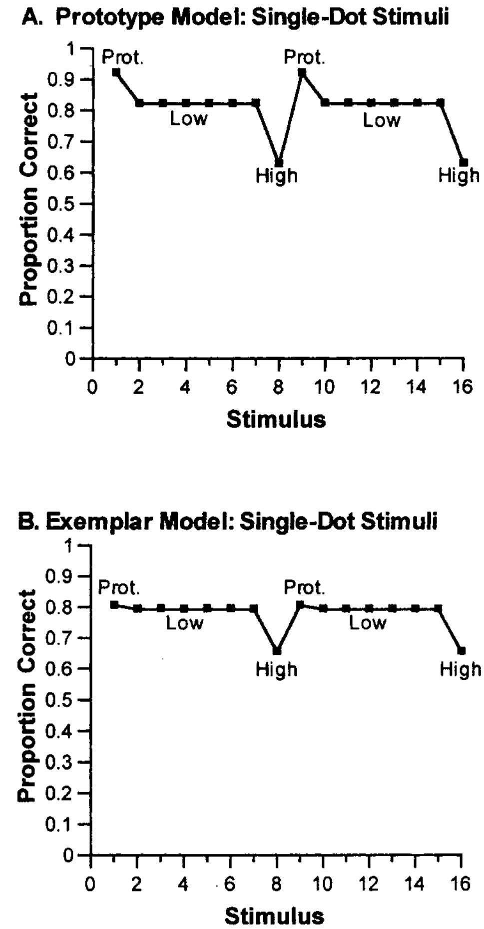

Figure 6. Panel A: The average performances predicted by a prototype model for stimuli at the category's center (Stimuli 1 and 9), near the category's center (Stimuli 2–7, 10–15), and far from the category's center (Stimuli 8 and 16). Panel B: The same for a comparable exemplar model. Prot. = prototype.

homogeneous performance profiles, flat typicality gradients, and small prototype advantages. The prototype model has the opposite performance characteristics.

Second, we defined two nine-dot prototypes and built categories containing the prototype of each category (Stimuli 1 and 9), six items that were 5-bits-per-dot distortions of each prototype (Stimuli 2–7, 10–15), and one item that was an 8.6-bits-per-dot distortion of each prototype (Stimuli 8 and 16). Then we chose 5,000 random sensitivities and found the predicted performance levels for these 16 stimuli according to the multiplicative prototype and exemplar models. The 5,000 performance profiles for each model were averaged into the graphs shown in Figures 7A and 7B. The different quality of prototype-based and exemplar-based performances is clear again.

Interpreting the Categorization Performance of Controls and Amnesics

These differences between prototype-based and exemplar-based performances can be used in interpreting the categorization per-

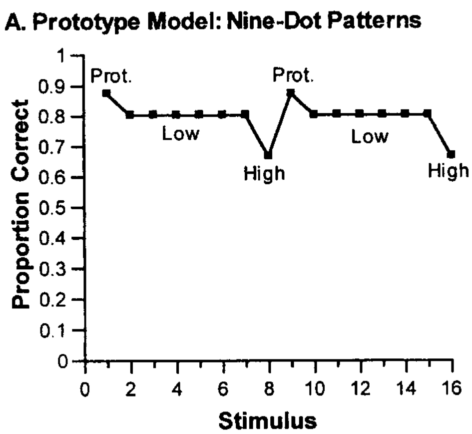

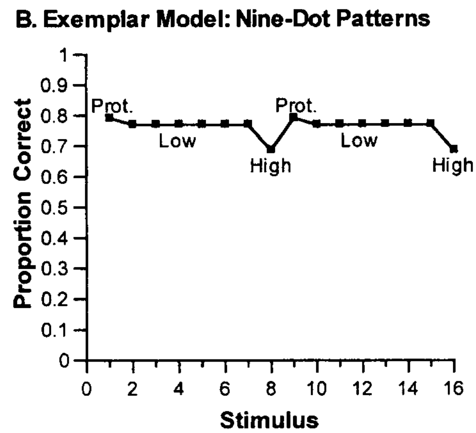

Figure 7. Panel A: The average performances predicted by a prototype model for prototype items (Stimuli 1 and 9), low-level distortions (Stimuli 2–7, 10–15), and high-level distortions (Stimuli 8 and 16). Panel B: The same for a comparable exemplar model. Prot. = prototype.

formances of Knowlton and Squire's (1993) controls and amnesics. Figures 8A and 8B show these performances for the four item types (prototypes, low-level distortion category members, high-level distortion category members, and random patterns outside the category), together with the performances predicted by the power model for each group (see also Table 1).

The power model seems to have somewhat the wrong characteristics for fitting these performances. Regarding small prototype advantages, it predicted only half that observed (6% vs. 12%). Regarding flat typicality gradients, it predicted less than half the observed change in performance across prototypes, low-level distortions, and high-level distortions (8% vs. 17%). Regarding homogeneous performance profiles, it predicted less than one fourth the observed variance across these three category-member item types (.001 vs. .005).

One possibility is that the power model has trouble with these three item types because they all lie on or inside the shell of high-level training distortions—in the region of category space where exemplar-based distance changes slowly relative to prototype-based distance. These three item types will be about equally distant from the shell, will seem about equally similar to the training exemplars, and should be placed by an exemplar model about equally strongly in the category. This equality is a geometrical consequence of comparisons to an exemplar shell.

However, the power model does show some (half) of the required prototype advantages. This makes it seem as though flat typicality gradients are not an absolute property of comparisons to a shell of training exemplars. However, our simulations suggest that they are. Thus, important concerns arise about the subjective similarity ratings that Nosofsky and Zaki (1998) input into their power model and about the behavior of the power model itself.

Similarity Ratings and Models of Categorization

Nosofsky and Zaki (1998) used subjective similarity ratings collected after category learning to model category learning. This approach is unusual. Usually, objective measures of stimulus similarity are used, or possibly subjective similarity measures collected before or absent category learning (e.g., Lamberts, 1995; McKinley & Nosofsky, 1996; Nosofsky, 1986, 1987; Nosofsky, Palmeri, & McKinley, 1994; Posner, 1964; Posner et al., 1967;

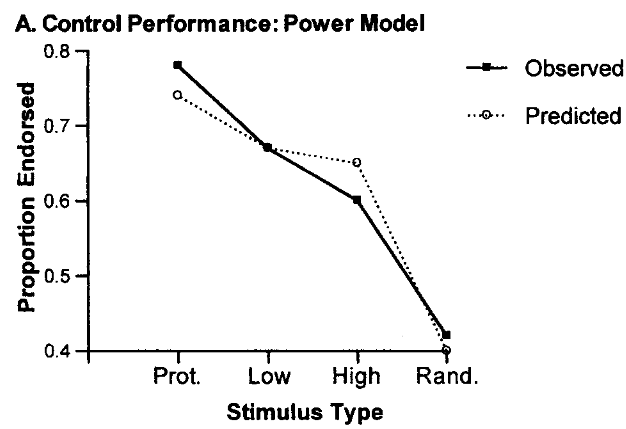

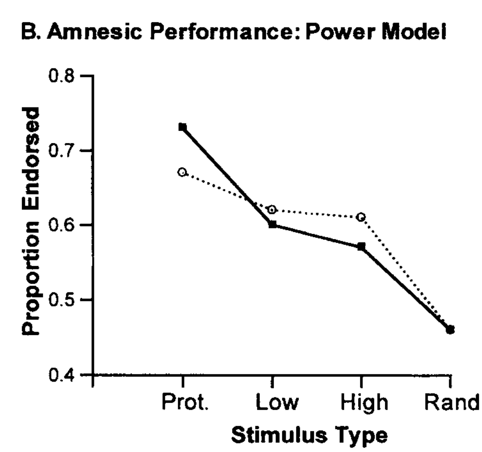

Figure 8. The fit of Nosofsky and Zaki's (1998) power model (dotted lines) to the observed categorization performances (solid lines) of control (A) and amnesic (B) participants in Knowlton and Squire's (1993) study. Prot. = prototype; Rand. = random.

Table 1
Observed and Predicted Performance Levels for Prototypes (P), Low-Level Distortions (L), High-Level Distortions (H), and Random or Unrelated Patterns (R)

Group and model	P	L	H	R	Fit measure
Knowlton and Squire (1993)					
Controls					
Observed	.78	.67	.60	.42	
Power	.74	.67	.65	.40	.0045
Exemplar (distance)	.70	.69	.65	.42	.0088
Prototype (distance)	.79	.67	.58	.43	.0008
Fine-grained exemplar	.71	.70	.65	.41	.0085
Gamma	.71	.70	.65	.41	.0086
Knowlton and Squire (1993)					
Amnesics					
Observed	.73	.60	.57	.46	
Power	.67	.62	.61	.46	.0056
Exemplar (distance)	.65	.64	.61	.45	.0098
Prototype (distance)	.72	.62	.56	.46	.0007
Fine-grained exemplar	.65	.65	.61	.45	.0098
Gamma	.65	.64	.61	.45	.0098
Reber, Stark, and Squire (1998b)					
Controls					
Observed	.71	.55	.54	.40	
Exemplar (distance)	.62	.61	.57	.40	.0129
Prototype (distance)	.70	.58	.51	.40	.0020
Reber, Stark, and Squire (1998a)					
Controls					
Observed	.85	.66	.63	.40	
Exemplar (distance)	.73	.72	.67	.40	.0193
Prototype (distance)	.84	.70	.59	.41	.0030
Knowlton and Squire (1993, Exp. 2)					
Controls					
Observed	.76	.71	.57	.31	
Exemplar (distance)	.70	.69	.63	.32	.0071
Prototype (distance)	.82	.66	.54	.35	.0089
Palmeri and Flanery (1999)					
No training					
Observed	.71	.61	.51	.37	
Fine-grained exemplar	.67	.62	.56	.35	.0038
Prototype (distance)	.72	.59	.51	.38	.0008
Most recent four items	.66	.61	.58	.47	.0165
Most recent four members	.69	.63	.57	.38	.0039

Note. Exp. = experiment.

Smith & Minda, 1998; Smith & Minda, 2000; Smith, Murray, & Minda, 1997).

One reason their approach is rarely adopted is that it could be circular. Similarity ratings among dot-pattern stimuli are affected by category training and reflective of category training. For example, Homa, Rhoads, and Chambliss (1979) showed that similarity ratings taken after category training reflected the coalescence of families of exemplars around their respective prototypes. Thus, one can see how using postlearning similarity ratings to model learning could compromise psychological interpretation. Suppose that participants learn the prototype during category learning and similarity space organizes around it. Then the prototype would have a privileged, higher similarity to the training exemplars after category training that similarity ratings would reflect. The power model would incorporate that higher similarity through the ratings and use it to predict higher prototype performance. But this would not imply any exemplar-based process at work. To the contrary, it would really confirm that prototype abstraction had been at work. The "exemplar" model that received these prototype-privileged similarity ratings as inputs would thus become psychologically empty or theoretically misleading in an important way.

In fact, one can show that Nosofsky and Zaki's (1998) postlearning similarity ratings, used as inputs to the power model, do cause this problem. Figure 9 shows the similarity ratings that Nosofsky and Zaki input into the power model, and the performance predictions the model output (for the prototype, the low-level distortions, the high-level distortions, and the random items) when it fit the data for Knowlton and Squire's (1993) control participants. The power model's predicted-performance outputs are just a simple arithmetic transformation of the similarity-rating inputs. The power model does not explain the inputs; it only reflects them. Thus, a real question arises as to whether the power model contains an exemplar process within it and as to whether it should be considered an exemplar-based model.

Therefore, theoretical interest steps back to focus on the postlearning similarity ratings themselves. Nosofsky and Zaki's (1998) participants could not have produced the pattern of similarity ratings they did (i.e., the strong boost in similarity ratings for the

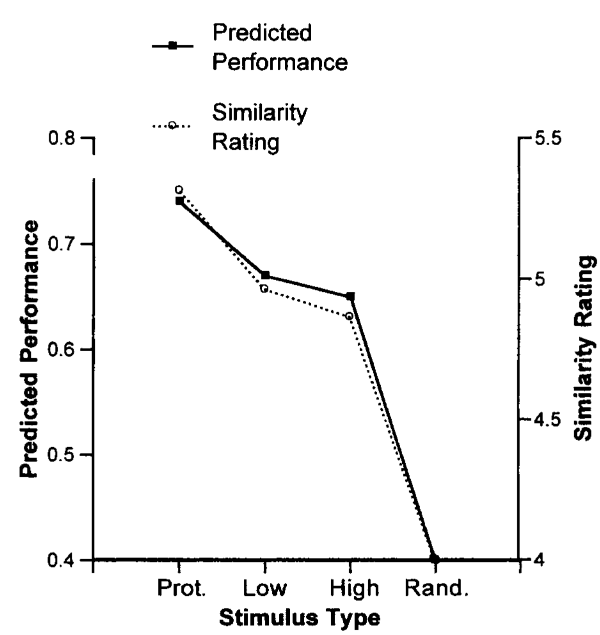

Figure 9. Mean rated similarity (dotted line, right vertical axis) between four test-item stimulus types (prototype, low-level distortion, high-level distortion, and random) and high-level distortion training items. These ratings were produced by Nosofsky and Zaki's (1998) participants following category learning. This figure also shows the levels of category endorsement (solid line, left vertical axis) predicted by the power model when in its best fitting configuration for fitting the data from Knowlton and Squire's (1993) control participants. Prot. = prototype; Rand. = random.

prototype) if they were actually making exemplar-to-exemplar comparisons. Posner et al. (1967) showed instead that exemplar-to-exemplar comparisons produce extremely flat similarity gradients moving in toward and including the prototype. That is, the prototype was no more similar to the training distortions than were low-level distortions or even high-level distortions. Posner et al.'s (1967) Table 3 (p. 34) alone showed 20 confirmations of this phenomenon (see also Figure 1C in the present article). We re-created this result (see Figure 2C in the present article). We also explained the intuitive reasons for this phenomenon and showed its robustness over different attentional strategies and similarity metrics.

If not exemplar-to-exemplar comparisons, then what processes are suggested by the strong changes in similarity ratings observed by Nosofsky and Zaki (1998)? We believe these strong changes reflect that category training has taught participants the prototype as the privileged organizing representation of the category's exemplars. In this way, when given comparisons between training distortions and the prototype, they would note the latter's ideal resonance as the originating pattern and grant it a special, higher similarity. This would explain why Nosofsky and Zaki (1998) observed a pattern different from that of Posner et al. (1967) and from our re-creation of their work. This would explain why it is category training that produces this different pattern. This would even explain why, for three of the four item types (random, high-level distortions, low-level distortions), the similarity gradient does flatten moving in toward the prototype and why it is only the prototype similarity that spikes upward again (see Figure 9, dotted line).

This interpretation has important implications. It suggests that the similarity ratings collected by Nosofsky and Zaki (1998) themselves reflect prototype abstraction. It means that the use of these ratings within the power model is potentially misleading psychologically and theoretically. For by doing so, one is taking prototype-privileged ratings, giving them a simple mathematical transform, miscalling that transform an exemplar process, and concluding that exemplar-based processing explains the data. Clearly this is a problematic procedure.

Therefore, we adopted instead a more neutral learning-free way to measure the psychological similarities among stimulus classes in the category task of Knowlton and Squire (1993) and to provide the inputs to formal models of categorization. As an alternative similarity measure, we chose the measure that was the fundamental result of Posner et al.'s (1967) similarity-scaling research with dot patterns. Thus we adopted the measure ln (1 + average Pythagorean distance moved by corresponding dots between patterns). Posner et al. (1967) demonstrated that this measure tracks perfectly the similarity ratings neutral observers give pairs of dot patterns. The experiment reported earlier in this article reached the identical conclusion. This measure has the additional advantages that it can be calculated objectively, that it is theoretically neutral regarding prototype and exemplar representations in categorization, and that it is calculated identically for prototype–exemplar and exemplar–exemplar pairs. This means that one can derive equivalent inputs to exemplar and prototype models and then ask how well the characteristics of equivalent exemplar and prototype models suit the characteristics of category learners.

The Behavior of the Exemplar Model Given Neutral Psychological Distances

Accordingly, we fit the categorization data of the control and amnesic groups from Knowlton and Squire (1993) using a distance-based exemplar model with the choice rule:

$$P(R_A|S_i) = \frac{\eta_{ih}}{\eta_{ih} + k}. \tag{8}$$

This equation is analogous to that for the power model because it compares total category similarity (η) with a criterion value (k) that is a free parameter in the model. As in the power model, this model used the average similarity between the to-be-categorized item type (i) and the high-level training distortions (h) as the estimate of category similarity.

However, this model calculated similarity as follows. We started with the average objective logarithmic distances of each transfer item type from the high-distortion training exemplars. We calculated these distances using the exact position of every dot in Knowlton and Squire's (1993) primary group of stimulus patterns.[1] We scaled each distance by a freely estimated sensitivity parameter as in the standard exemplar model. (Thus, this distance-based exemplar model has two free parameters—sensitivity and criterion.) Then we transformed this scaled distance into a similarity using an exponential-decay function.

Given estimates of similarity, we found the predicted level of category endorsement for each transfer item type (prototypes, low-level distortions, high-level distortions, and random items) by comparing similarity with the criterion quantity, k. We used a hill-climbing algorithm to maximize the fit between predicted and observed profiles and thus to find the best-fitting predicted profile.

Figure 10 shows the two observed performances with the best-fitting predictions of the distance-based exemplar model (see also Table 1). Given a neutral measure of psychological similarity, the exemplar model's best fitting predicted performance profiles are qualitatively different from the observed performance profiles. The exemplar model predicts almost none of the observed prototype advantage and almost no typicality gradient over the three category-member stimulus classes.

The exemplar model fits poorly because the three category-member stimulus types—prototypes, low-level distortions, and high-level distortions—lie on or within the shell of training exemplars and have about the same distance to the shell. Concretely, the Pythagorean distance from training exemplars to the prototype, the low-level distortions, and the high-level distortions was 4.52, 4.80, and 5.97, respectively. The corresponding logarithmic distances were 1.71, 1.76, and 1.94, respectively. These equivalent distances give rise to equivalent similarities, equivalent predicted response proportions, flat typicality gradients, and small prototype advantages. In short, the performances predicted by the exemplar model have exactly the properties that are diagnostic of similarity comparisons to an exemplar shell. But this is not what participants show.

---

[1] We thank Safa Zaki for providing these coordinates to us.

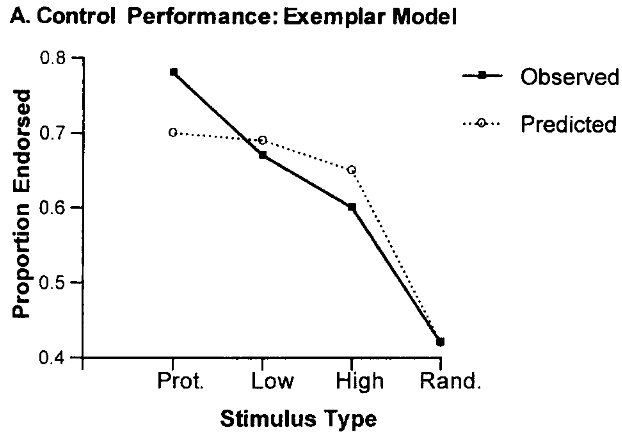

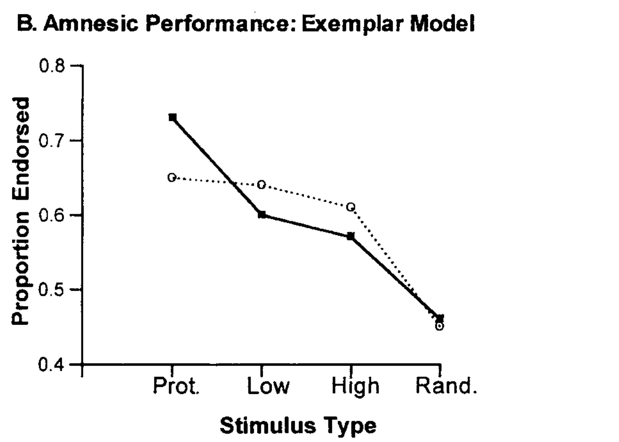

Figure 10. The fit of a distance-based exemplar model (dotted lines) to the observed categorization performances (solid lines) of control (A) and amnesic (B) participants in Knowlton and Squire's (1993) study. Prot. = prototype; Rand. = random.

The Behavior of a Matched Prototype Model Given Neutral Psychological Distances

We also fit both data sets using a distance-based prototype model that was identical in every respect to the exemplar model just discussed, except in its assumption of a prototype category representation. The choice rule for this model was

$$P(R_A|S_i) = \frac{\eta_{iP}}{\eta_{iP} + k}. \tag{9}$$

This model calculated the total category similarity (η) between a transfer item type (i) and the prototype (P) by starting with the average objective logarithmic distance between tokens of the item type and the prototype. These distances were also found with Knowlton and Squire's (1993) stimulus specifications. Each distance was scaled by a freely estimated sensitivity parameter, converted into a similarity using an exponential-decay function, then converted into a predicted level of category endorsement through comparison with the freely estimated criterion quantity, k. Again we used a hill-climbing algorithm to maximize the fit between predicted and observed performance profiles and thus to find the best fitting predicted profile.

Figure 11 shows the two observed performances with the best fitting predictions of the distance-based prototype model (see also Table 1). The prototype model fits the data well. Its fitting error is one tenth that of the equivalent exemplar model. The prototype model fits performance well because the distances from the prototype to all four item types—prototypes, low-level distortions, high-level distortions, and random items—increase steadily and strongly. Concretely, the Pythagorean distance from the prototype to the prototype, the low-level distortions, the high-level distortions, and the random items was 0.00, 2.08, 4.82, and 15.07, respectively. The corresponding logarithmic distances were 0.00, 1.13, 1.76, and 2.77, respectively.[2] Rated similarity and predicted levels of categorization would follow this logarithmic distance and decrease strongly across the item types, producing large prototype advantages and strong typicality effects. In short, the performances predicted by the prototype model have the characteristics that are diagnostic of similarity comparisons to a category center, and this is what participants show.

An Alternative Exemplar Model Using Individual-Item Distances and Similarities

We also evaluated an exemplar model that used the fine-grained, item-by-item analyses that are truest to the analytic spirit of exemplar theory. We thought this approach might give the exemplar model a better chance of succeeding. The choice rule for this model was

$$P(R_A|S_i) = \frac{\sum_{h \in C_A} \eta_{ih}}{\sum_{h \in C_A} \eta_{ih} + k_{40}}. \tag{10}$$

Here we referred each of the 84 transfer items (i) back to each of the 40 high-distortion training exemplars (h), found each of the 40 logarithmic distances individually, scaled each distance by sensitivity (one free parameter in the model), and turned each scaled distance into a similarity (η) using an exponential-decay function. Then we summed the 40 similarities each transfer item had to the 40 category members and found its predicted level of category endorsement by comparing this summed similarity with the criterion quantity (k_{40}—the other free parameter in the model).

[2] In calculating these distances for the random items, the potential problem arises of which dots correspond between the random item and the prototype. To evaluate this problem, we compared the 40 random items to the prototype using 100 different dot-to-dot correspondences for each pattern. The 100 estimates of the Pythagorean distance from the prototype to the 40 random transfer items averaged 14.86 (SD = 0.29). The 100 estimates of the logarithmic distance averaged 2.76 (SD = 0.02). The distribution of logarithmic distance estimates is so narrow that any value from it would produce nearly identical modeling results. This means that ambiguity about dot correspondences has no impact on the results reported here. In addition, one should remember that all models agree on the extreme distance of the random items from the category members and on the low levels of category endorsement predicted for them. The issues of theoretical importance concern stimuli close in to the prototype where there is no ambiguity about dot correspondences.

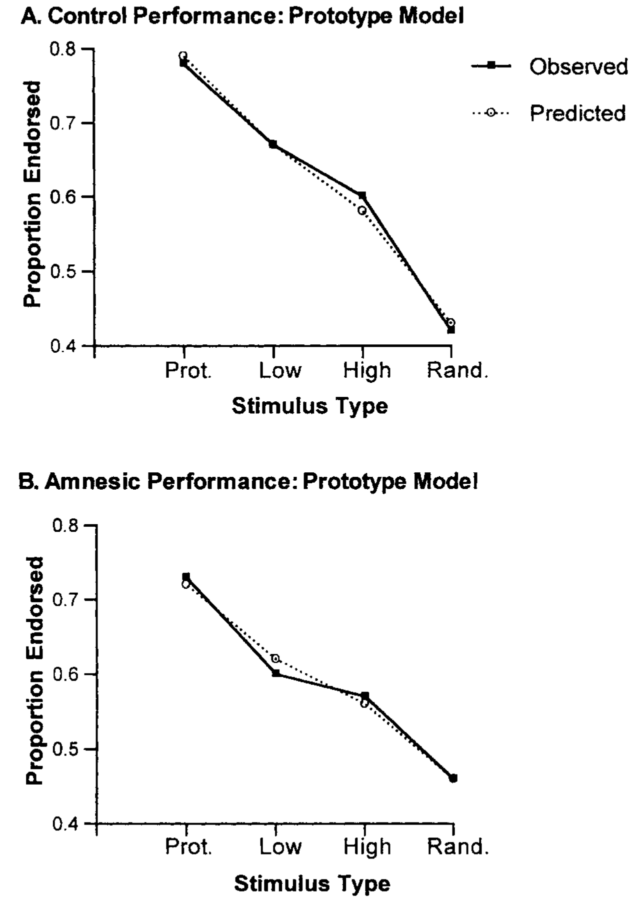

Figure 11. The fit of a distance-based prototype model (dotted lines) to the observed categorization performances (solid lines) of control (A) and amnesic (B) participants in Knowlton and Squire's (1993) study. Prot. = prototype; Rand. = random.

This process was repeated for each of the 84 transfer items, and then these were collected by averaging into transfer item types to determine the level of category endorsement that a configuration of the model predicted for prototypes, low-level distortions, high-level distortions, and random items. These four predicted values were compared with the four observed values, and we used a hill-climbing algorithm to maximize the fit between the predicted and observed values and thus to find the best fitting predicted profile. This fine-grained exemplar model produced best fitting predictions that were almost identical to the average exemplar model (see Table 1), and so this approach did not let the exemplar model fit better.

An Alternative Exemplar Model Incorporating Gamma

We also evaluated an exemplar model that incorporated the additional parameter gamma. This parameter has sometimes improved substantially the fit of the exemplar model, so we thought this approach might give the exemplar model its best chance of succeeding. The choice rule equation for this model was

$$P(R_A|S_i) = \frac{\eta_{ih}{}^{\gamma}}{\eta_{ih}{}^{\lambda}+k^{\gamma}}. \qquad (11)$$

This model calculated the total category similarity (η) between a transfer item type (i) and the high-level training distortions (h) by starting with the average objective logarithmic distance between tokens of the item type and the training exemplars. Each distance was scaled by a freely estimated sensitivity parameter, converted into a similarity using an exponential-decay function, then converted into a predicted level of category endorsement through comparison with the freely estimated criterion quantity, k. In the final choice-rule equation, though, both the quantities (category similarity of the item type in the numerator and denominator, and criterion quantity in the denominator) could be raised to any power (this model's third free parameter γ) that let the model best recover the four observed categorization performances. We used a hill-climbing algorithm to maximize the fit between predicted and observed profiles and find the best fitting predicted profile. The exemplar model with gamma produced best fitting predictions that were almost identical to the exemplar model without gamma (see Table 1), and so this approach did not help the exemplar model fit better.

We pause to point out that the contrast between the geometries of exemplar-based and prototype-based comparisons explains a concern about including gamma in the exemplar model. Gamma lets the exemplar model estimate high performances higher and low performances lower. It counters the inherent tendency for exemplar-based comparisons to produce flat typicality gradients. It can formally steepen typicality gradients in a way that mimics a prototype-centered geometry. Smith and Minda (1998) illustrated this using extensive simulations. For this reason, gamma—when added to the exemplar model—has potential representational implications that could make the exemplar model less theoretically clear and possibly less exemplar-based. This is why Smith and Minda (1998) concluded that the inclusion of the gamma parameter in exemplar models needs careful consideration.

Additional Data Sets

Our principal focus in this article is on the data set (Knowlton & Squire, 1993) that Nosofsky and Zaki (1998) highlighted in their interpretation of the amnesia dissociation. Here we briefly note additional applications of a perspective that focuses on the character of typicality gradients.

Reber et al. (1998a, 1999b). One can examine the character of typicality gradients in any data set to judge whether it indicates prototype-based processing. Two recent positive examples can be found in Reber et al.'s work (1998a, Figure 2; 1998b, Figure 2). In both experiments, participants studied 40 high-level distortions of a prototype and then were tested on that training prototype (4 tokens), low-level distortions of it (20 or 16 tokens in the two experiments, respectively), new high-level distortions of it (20 or 16 tokens), and items that were unrelated to it (40 or 36 tokens). In one case (Reber et al., 1998b), the unrelated items were high-level distortions of different randomly chosen prototypes. In the other case (Reber et al., 1998a), the unrelated items comprised 4 tokens of an untrained prototype, 16 of its low-level distortions, and 16 of its high-level distortions.

Table 1 summarizes the observed data and the results of fitting equivalent distance-based exemplar and prototype models to them. In both cases, the exemplar model predicts typicality gradients that are too flat to capture the real data, so that its fit index is much larger than the prototype model's fit index. In both cases, the strong typicality effects observed suggest prototype-based processing more than exemplar-based processing.

Knowlton and Squire (1993), Experiment 2. Second, we consider an interesting negative example to show that the character of typicality gradients can also indicate poor or incomplete abstraction under theoretically appropriate circumstances. To equate the repetition of information and the memory requirements across the categorization and recognition tasks, Knowlton and Squire (1993) also tested categorization after participants had studied only four training patterns. The literature strongly suggests that one should expect to see poor or incomplete abstraction under these conditions (Brooks, 1978; Homa et al., 1981; Smith & Minda, 1998; Smith et al., 1997) because participants have not been given an exemplar burden that requires abstraction or broad enough exemplar information to ensure its success. Table 1 summarizes the observed data from control participants and the results of fitting the distance-based exemplar and prototype models to them. Here the typicality gradients are flatter. Here the character of performance does not strongly suggest prototype-based processing, and the fit index for the prototype model is slightly higher than that for the exemplar model.

Palmeri and Flanery (1999). Finally, we ask whether the character of typicality gradients can illuminate an intriguing recent demonstration of Palmeri and Flanery (1999). They replicated Knowlton and Squire's (1993) experiment except that they actually presented participants with no training patterns at all (though participants were told they had seen the training patterns subliminally). These participants were able to categorize the transfer patterns at about the levels seen in Knowlton and Squire (1993). Apparently, participants quickly inferred that the category members were the set of highly similar patterns in the transfer series, whereas the nonmembers were the random, dissimilar patterns. This inference is valid because the transfer set was composed of similar distortions of a prototype (to be judged members) and a set of completely random patterns (to be judged nonmembers). This result provides a useful note of caution into the discussion of exactly what controls and amnesics learn in training; one possibility is that they learn little but infer smartly during transfer. In fact, Palmeri and Flanery noted that this result embodies a potential criticism of both traditional exemplar-based and prototype-based accounts of categorization, because both typically view the training phase as the time when exemplars are stored or prototypes abstracted.

Given the interesting result that the cognitive work of category learning can occur during the transfer phase, it is worth asking what participants learn then. In accordance, we fit the distance-based prototype and exemplar models to Palmeri and Flanery's (1999) data. In this case we used the fine-grained exemplar model previously described and granted the exemplar model 44 exemplar traces (4 prototypes, 20 low-level distortions, and 20 high-level distortions), consistent with the possibility that Palmeri and Flanery's participants were storing the transfer exemplars and using them to categorize. This approach gives the exemplar model a strong assist in reproducing sharp typicality gradients, because the stored low-level distortions (20) and stored prototypes (4) help it predict high levels of performance for prototypes and low-level distortions. Table 1 summarizes the observed data and the results of fitting the exemplar and prototype models. The exemplar model predicts too flat a typicality gradient to capture the real data, so that its fit index is larger than that of the prototype model. So the strong typicality gradients observed suggest that Palmeri and Flanery's participants used something more like prototype-based processing than like exemplar-based processing.

We also modeled the Palmeri and Flanery (1999) data set in two additional ways. In one case, we assumed that each test exemplar was only compared to several of the stimuli that had just been seen. Thus, if an item shared a lot of similarity with the four most recently presented items (which would happen if the item and the four most recently presented items were category members), it could reasonably be endorsed as a member of the category. But if it did not share similarity (which would happen if the item and the four most recently presented items were not category members), it should not be so endorsed.

We instantiated these assumptions into a model that could be fit to the data in the following way. We first picked a level of sensitivity and a criterion (as in the previous model fitting). For each of the 84 transfer stimuli, we then randomly chose four exemplars from among the 84 to be its comparison standard. Using the fine-grained exemplar model already discussed, a prediction was made for the performance on the stimulus. We repeated the choice of four random comparison exemplars 100 times for each stimulus, to estimate comprehensively how strongly that item would be endorsed across all contexts of four recently seen items. This process of choosing 100 sets of four comparison exemplars was repeated for each of the 84 transfer items, and then performance levels were collected into transfer item types to determine the overall level of category endorsement that a configuration of the model gave prototypes, low-level distortions, high-level distortions, and random items. These four predicted values were compared with the observed values, and we used a hill-climbing algorithm to maximize the fit between the predicted and observed values and find the best fitting predicted profile. Table 1 summarizes the results of fitting this model to the data. This approach did not help the exemplar model fit better.

In another case, we repeated the process just described exactly, except that we assumed instead that participants were holding on to just the four most recently seen actual members of the category. That is, the four comparison exemplars were always chosen from among the 44 transfer items that were members of the category instead of being chosen from among the whole set of 84 transfer items. Table 1 summarizes the results of fitting this model to the data. This approach did not help the exemplar model fit better.

In summary, a careful analysis of typicality gradients also indicates prototype-based processes in other recent data sets. It illuminates the findings of Palmeri and Flanery (1999). It points to cases of poor or incomplete abstraction just where these might be expected. Thus, an approach that examines carefully the shape of typicality gradients offers the promise of illuminating the kinds of category representations and comparison processes that produced them.

Summary of Theoretical Analysis: Categorization

The exemplar-based power model of Nosofsky and Zaki (1998) fit the categorization performances of controls and amnesics fairly poorly, and only that well because its postlearning similarity-rating inputs reflected the prototypes acquired during category learning.

When neutral similarity measures (measures grounded in extensive similarity-scaling studies) were used instead, the exemplar model fit the observed performances very poorly. Alternative exemplar models that used more fine-grained exemplar-based information or that incorporated the gamma parameter fit poorly as well.

Exemplar models fail to fit these data because of definite geometrical laws that follow from comparing to-be-categorized items to the shell of exemplars that surround a prototype. Exemplar-based distance, similarity, and performance level all change slowly close to the prototype, flattening typicality gradients, shrinking prototype advantages, and homogenizing performance profiles. As a result, exemplar models do not account for the strong typicality effects in the observed results.

In contrast, a prototype model given the same neutral, psychological distance inputs fits these data well. Prototype models fit these data because of definite geometrical laws that follow from comparing to-be-categorized items to the center (the prototype) of the category. Prototype-based distance, similarity, and performance level still can change strongly close to the prototype, keeping typicality gradients sharp, prototype advantages large, and performance profiles heterogeneous.

As to the observed data, the strong typicality gradients, the large prototype advantages, and the heterogeneous performance profiles suggest that participants are probably referring to-be-categorized items to the center of the category in making their category decisions. If so, this analysis establishes that one memory process underlying the categorization-recognition dissociation is prototype based, recalling the traditional interpretation. In turn, this produces the clear expectation that explaining the categorization-recognition dissociation will require multiple memory processes, because we can confidently expect the recognition task, by its nature, to use some exemplar-based process.

Recognition

The possible confirmation of a prototype-based process underlying categorization helps as one tries to understand the exemplar process underlying recognition. It allows one to reconsider the idea from exemplar theory that recognition is based on systematic similarity comparisons of an item back to all the stored, individuated exemplar traces from the recognition list (Nosofsky, 1988, 1991; Palmeri & Nosofsky, 1995). If the categorization performances observed by Knowlton and Squire (1993) were not based on these systematic exemplar comparisons, the recognition performances they observed need not have been either. Instead, it is possible to consider different and simpler exemplar-based processes. This simplification could link the present article to recent articles on categorization. For example, Smith and Minda (2000) showed that combining prototype-based processing with single-comparison exemplar memorization often accounts well even for categorization performances that have normally been analyzed using the standard exemplar model with its multiple and systematic exemplar-to-exemplar comparisons. We ask whether recognition is another case in which exemplar memorization or exemplar self-retrieval suffices. If so, we could let this simple exemplar process on the recognition side complement the prototype process on the categorization side, and explain the categorization-recognition dissociation in a way that reflects recent findings in categorization and past interpretations of the amnesia dissociation. This possibility is the topic of the next sections of the article.

Self-Similarity in Old-Item Recognition

The key assumption of current exemplar-based models of recognition is that a to-be-recognized item is compared to all the members of the recognition list and is recognized if its total similarity to those items exceeds some criterion amount. In accordance, Nosofsky and Zaki's (1998) equation for old-item recognition in their power model is

$$P(R_{old}|S_i) = \frac{SelfSim^p + 4 \times sr(i, r)^p}{SelfSim^p + 4 \times sr(i, r)^p + k_5}. \quad (12)$$

The numerator summarizes the total similarity a to-be-recognized item has to the five members of the recognition list. Two sorts of similarity apply. First, the item is compared to its own exemplar trace in memory (*SelfSim*). The *SelfSim* quantity was estimated as a free parameter in the model (its best fitting value was 9.72). Second, the item is compared to the four other random unrelated items in the recognition list. To calculate this quantity ($4 \times sr(i, r)$), Nosofsky and Zaki used the average similarity rating that participants gave for pairs of random patterns. Both similarity quantities were also raised to a power to preserve the idea in the power model that true psychological similarity is a power function of the similarity ratings participants actually offer. This parameter let the model account for different recognition sensitivity by controls and amnesics. Finally, k_5 is the criterion quantity or proportionalizing threshold. Intuitively, one can take this quantity as the total similarity of the to-be-recognized item to five null or baseline items. The more the total similarity to the 5 actual members of the recognition list exceeds this baseline total, the more strongly it would be endorsed as old. In all, the power model uses three free parameters to fit recognition data (*SelfSim*, *p*, and k_5).

Thus, for old-item recognition, the probability of a recognition response equals the impact of the self-similarity of the item with its own trace in memory plus the impact of the four similarities between the item and the four other recognition-list items—all compared with the criterion for accepting the item as old. Note that this recognition model, but not the categorization model already discussed, uses self-similarity to treat specific items in an individuated, not averaged, way. This model lets old exemplars receive privileged processing. It emphasizes specific-exemplar processes, whereas the categorization model downplayed them. The availability of self-similarity in the recognition model shows that Nosofsky and Zaki's (1998) categorization and recognition models are somewhat different and potentially include different processes.

Control Participants and Old-Item Recognition

Moreover, the power model's best fitting quantitative description of control participants' old-item recognition shows that self-similarity plays a crucial role in the recognition model's behavior and success:

$$P(R) = \frac{(9.72^{5.18} + 4 \times 4.139^{5.18})}{(9.72^{5.18} + 4 \times 4.139^{5.18} + 28{,}877)} \quad (13)$$

or

$$P(R) = \frac{(130{,}650 + 4 \times 1{,}568)}{(130{,}650 + 4 \times 1{,}568 + 28{,}877)}. \tag{14}$$

The dominant feature of this equation is that exemplar self-retrieval makes an impact on the cognitive system of controls of 130,650. This means that a single event of exemplar self-retrieval has about as strong an impact as the presumed accumulated similarity from all 40 training items in a trial of the category task. This means that the impact from self-retrieval is 83 times larger than that from comparing the to-be-recognized item to other recognition-list members (130,650 vs. 1,568). The relative weakness of the second exemplar process could be expected because the members of the recognition list are very distant from each other (18.34 Pythagorean units per dot) and minimally similar to each other. But this means that the second exemplar process might be negligible in controls, leaving self-recognition alone to govern decision making in old-item recognition.

To evaluate this possibility, we eliminated the second exemplar process from the power model's recognition equation. That is, we rewrote

$$P(R) = \frac{(130{,}650 + 4 \times 1{,}568)}{(130{,}650 + 4 \times 1{,}568 + 28{,}877)} \tag{15}$$

as

$$P(R) = \frac{(130{,}650)}{(130{,}650 + 28{,}877)}. \tag{16}$$

The original equation (Equation 15) that assumes two exemplar processes in recognition predicts that old items will be recognized 82.6% of the time. Equation 16 that assumes only exemplar self-retrieval predicts that old items will be recognized 81.9% of the time. The second exemplar process is negligible. The situation is like reading at night in the light cast by one movie-premier spotlight and four matches. You would read by the spotlight. The more complex description of the light is unparsimonious.

Alternatively, consider a thought experiment. If one shows picture postcards of the Parthenon, Cheop's Pyramid, the Washington Monument, the White House, and the Great Wall of China and later tests the Parthenon for recognition, probably the Parthenon's resonance with itself in memory alone ensures recognition. It seems less likely that the recognition decision is supported by other similarity comparisons (e.g., the Parthenon–Washington Monument similarity), especially after this second comparison process is eliminated from the equation without affecting predicted recognition at all.

Amnesic Participants and Old-Item Recognition

Regarding old-item recognition by amnesics, the quantitative description of the power model suggests that exemplar self-recognition makes an impact 145 times less on the cognitive systems of amnesics ($9.72^{2.99}$ or 897) than on the cognitive systems of controls ($9.72^{5.18}$ or 130,650). The difference confirms that controls use self-retrieval more powerfully in the recognition task than do amnesics. It is still unclear whether secondary exemplar-comparison processes to other recognition-list members are required.

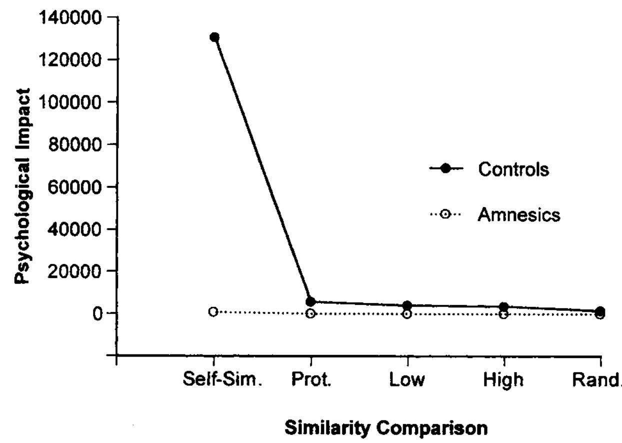

Figure 12. The psychological impact on controls and amnesics estimated by the power model when old exemplars self-retrieve within the recognition task, or when prototypes, low-level distortion category members, high-level distortion category members, and random nonmembers of the category are compared to a training exemplar within the category task. Prot. = prototype; Rand. = random; Self-Sim. = self-similarity.

Figure 12 summarizes the situation regarding the recognition and categorization performance of controls and amnesics from the perspective of the power model. It shows the impact that the power model estimates is made on the cognitive systems of controls and amnesics by different kinds of similarity comparisons (the comparison in the recognition task between an old item and its own trace in memory, and the comparison in the category task between training items and prototypes, low-level distortions, high-level distortions, and random items). The psychological landscape is dominated by the impact of self-recognition on the cognitive processing of controls. All other similarity comparisons make far smaller impacts on cognitive processing. Figure 12 makes clear that there is a big difference between categorization and recognition tasks. The difference is that the process of exemplar memorization or self-recognition operates powerfully in the recognition task but not in the categorization task. Figure 12 makes clear that there is a big difference between controls and amnesics. The difference is that the exemplar self-recognition process operates powerfully for controls but not for amnesics. This summary—that recognition is about exemplar memorization but categorization is not, and that controls use exemplar memorization but amnesics do not—fits exactly with Knowlton and Squire's (1993) original interpretation of the amnesia dissociation.

Smith and Minda (2000) also found that exemplar self-retrieval can dominate and supercede exemplar theory's more systematic exemplar-to-exemplar comparisons. Further, they showed the explanatory power of combining single-comparison exemplar memorization with a prototype-based process. If the recognition performance of controls and amnesics shows their differential use of the process of exemplar memorization, and if this process combines with a prototype-based process pointed to earlier in this article, then the categorization–recognition dissociation in amnesia is explained in a way that reinforces the conclusions of Smith and Minda and recalls the conclusions of Knowlton and Squire (1993) and others. This explanation could even extend to the extreme categorization–recognition dissociation shown by E.P. (Squire &

Knowlton, 1995; Squire & Zola, 1996) that Nosofsky and Zaki (1998) acknowledged the exemplar-based power model could not explain.

Conclusion

A crucial point of this article is that exemplar models operate according to definite geometrical laws that follow from comparing to-be-categorized items to the shell of exemplars that surrounds a prototype. Exemplar-based distance, similarity, and predicted performance levels all change slowly close to the prototype, flattening typicality gradients, shrinking prototype advantages, and homogenizing performance profiles. In contrast, prototype models operate according to definite geometrical laws that follow from comparing to-be-categorized items to the center of the category. Prototype-based distance, similarity, and performance still change strongly close to the prototype, keeping typicality gradients sharp, prototype advantages large, and performance profiles heterogeneous. The difference between these two descriptions of performance can let one think intuitively about categorization profiles and the representations and processes underlying them. For example, looking back now at the categorization performances observed by Knowlton and Squire (1993; Figure 11 in the present article), one can infer from the steep typicality gradients—even absent any modeling—that those performances were prototype based.

Several different exemplar models were evaluated to see whether they could fit these performance profiles. Granting an exemplar model systematic exemplar-to-exemplar comparisons or the gamma parameter did not help it fit better. Providing an exemplar model with similarity ratings obtained after category learning allowed it to predict only half the required prototype advantage and half the required typicality effect. Thus, these data represent a significant failure of exemplar theory.

In contrast, these data represent a significant success for prototype theory. This success joins other recent research showing the importance of prototype-based descriptions of individual participants' categorization performance (Smith et al., 1997), categorization in the early and middle stages of learning (Smith & Minda, 1998), performances in the category tasks that originally motivated exemplar theory (Minda & Smith, 2001a; Smith & Minda, 2000), and performance in a wide range of category tasks (Minda & Smith, 2001b).

Regarding the categorization–recognition dissociation in amnesia, the data suggest that there is a prototype-based process that serves categorization in controls and amnesics. This provisionally explains one half of the dissociation, endorses Knowlton and Squire's (1993) original interpretation, and reinstates the dominant interpretation of one of neuroscience's influential results.

Addressing the recognition half of the dissociation, this article showed that a single-comparison exemplar-memorization process was sufficient to explain the old-item recognition performance of controls and the weaker old-item recognition of amnesics. This interpretation is not opposed to the idea of exemplar representations or exemplar processes in cognition. To the contrary, we suggest that an exemplar process (self-retrieval) does underlie old-item recognition. This exemplar-memorization process provisionally explains the recognition half of the dissociation in the traditional way (i.e., good exemplar self-retrieval or strong exemplar self-resonance is good recognition). This explanation adds a second, exemplar-based process to the prototype-based process that serves categorization. Thus, it favors the important idea in the literature of multiple memory processes or systems underlying categorization and recognition.

Only one aspect of current exemplar theory is challenged by the analyses in this article. The systematic exemplar-to-exemplar comparison processes assumed by exemplar theory were not applicable to describing the categorization performances of Knowlton and Squire's (1993) controls and amnesics. (Indeed, the power model used by Nosofsky & Zaki, 1998, did not even allow for those exemplar-to-exemplar comparisons.) This inapplicability shifted the weight of parsimony and made it possible to see that these complex processes were not necessary for describing recognition performance either. Thus, there appears to be little support, regarding the performance of either group in either task, for exemplar theory's complex and systematic exemplar-to-exemplar comparisons. This conclusion—that neither categorization nor recognition performance is explained by those systematic exemplar comparisons—reinforces that of Smith and Minda (2000), who found that the standard exemplar model often fits categorization data by letting simple exemplar memorization dominate and by making negligible systematic exemplar-to-exemplar comparisons. We saw here in the recognition data just the same dominance of exemplar memorization and the same fading of the complex exemplar process.

The broader perspective of this article is that a single-comparison exemplar-memorization process joins a prototype-based process to explain the categorization–recognition dissociation in amnesia. This conclusion also reinforces recent research (Minda & Smith, 2001a, 2001b; Smith & Minda, 1998, 2000) and makes it less likely that a unified description of cognitive performance—one based in systematic exemplar-to-exemplar similarity comparisons—will be able to explain performances like categorization and recognition. Instead, the present article makes it seem more likely that a mixed description will be called for that is grounded in prototypes but includes exemplar memorization as well.

References

Anderson, J. R., Kline, P. J., & Beasley, C. M. (1979). A general learning theory and its application to schema abstraction. In G. H. Bower (Ed.), *The psychology of learning and motivation* (Vol. 9, pp. 269–315). New York: Academic Press.

Brooks, L. R. (1978). Non-analytic concept formation and memory for instances. In E. Rosch & B. B. Lloyd (Eds.), *Cognition and categorization* (pp. 169–211). Hillsdale, NJ: Erlbaum.

Hayes-Roth, B., & Hayes-Roth, F. (1977). Concept learning and the recognition and classification of exemplars. *Journal of Verbal Learning and Verbal Behavior, 16,* 321–338.

Hintzman, D. L. (1986). Schema abstraction in a multiple trace memory model. *Psychological Review, 93,* 411–421.

Hintzman, D. L., & Ludlam, G. (1980). Differential forgetting of prototypes and old instances: Simulation by an exemplar-based classification model. *Memory & Cognition, 8,* 378–382.

Homa, D., Cross, J., Cornell, D., Goldman, D., & Schwartz, S. (1973). Prototype abstraction and categorization of new instances as a function of number of instances defining the prototype. *Journal of Experimental Psychology, 101,* 116–122.

Homa, D., Dunbar, S., & Nohre, L. (1991). Instance frequency, categorization and the modulating effect of experience. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 17,* 444–458.

Homa, D., Rhoads, D., & Chambliss, D. (1979). Evolution of conceptual structure. *Journal of Experimental Psychology: Human Learning and Memory, 5,* 11–23.

Homa, D., Sterling, S., & Trepel, L. (1981). Limitations of exemplar-based generalization and the abstraction of categorical information. *Journal of Experimental Psychology: Human Learning and Memory, 7,* 418–439.

Knowlton, B. J., Mangels, J. A., & Squire, L. R. (1996, September 6). A neostriatal habit learning system in humans. *Science, 273,* 1399–1402.

Knowlton, B. J., & Squire, L. R. (1993, December 10). 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.

Lamberts, K. (1995). Categorization under time pressure. *Journal of Experimental Psychology: General, 124,* 161–180.

McKinley, S. C., & Nosofsky, R. M. (1996). Selective attention and the formation of linear decision boundaries. *Journal of Experimental Psychology: Human Perception and Performance, 22,* 294–317.

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

Minda, J. P., & Smith, J. D. (2001a). *Comparing prototype-based and exemplar-based accounts of category learning and attentional allocation.* Manuscript submitted for publication.

Minda, J. P., & Smith, J. D. (2001b). Prototypes in category learning: The effects of category size, category structure, and stimulus complexity. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 27,* 775–799.

Nosofsky, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General, 115,* 39–57.

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 categorization and recognition memory. *Journal of Experimental Psychology: Human Perception and Performance, 17,* 3–27.

Nosofsky, R. M., Palmeri, T. J., & McKinley, S. K. (1994). Rule-plus Exception Model of classification learning. *Psychological Review, 101,* 53–79.

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

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

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., & Nosofsky, R. M. (1995). Recognition memory for exceptions to the category rule. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21,* 548–568.

Posner, M. I. (1964). Uncertainty as a predictor of similarity in the study of generalization. *Journal of Experimental Psychology, 62,* 113–118.

Posner, M. I., Goldsmith, R., & Welton, K. E. (1967). Perceived distance and the classification of distorted patterns. *Journal of Experimental Psychology, 73,* 28–38.

Posner, M. I., & Keele, S. W. (1970). Retention of abstract ideas. *Journal of Experimental Psychology, 83,* 304–308.

Reber, P. J., Stark, C. E. L., & Squire, L. R. (1998a). Contrasting cortical activity associated with category memory and recognition memory. *Learning and Memory, 5,* 420–428.

Reber, P. J., Stark, C. E. L., & Squire, L. R. (1998b). Cortical areas supporting category learning identified using functional MRI. *Proceedings of the National Academy of Sciences, USA, 95,* 747–750.

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

Smith, J. D., & Minda, J. P. (1998). Prototypes in the mist: The early epochs of category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 24,* 1411–1430.

Smith, J. D., & Minda, J. P. (2000). Thirty categorization results in search of a model. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26,* 3–27.

Smith, J. D., Murray, M. J., & Minda, J. P. (1997). Straight talk about linear separability. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23,* 659–680.

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 nondeclarative memory systems. *Proceedings of the National Academy of Sciences, USA, 93,* 13515–13522.

Strange, W., Keeney, T., Kessel, F. S., & Jenkins, J. J. (1970). Abstraction over time of prototypes from distortions of random dot patterns: A replication. *Journal of Experimental Psychology, 83,* 508–510.

Received July 11, 2000
Revision received October 11, 2000
Accepted October 25, 2000 ■