

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Prototype Abstraction in Category Learning?

Permalink

<https://escholarship.org/uc/item/1q804664>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 23(23)

ISSN

1069-7977

Authors

Palmeri, Thomas J.
Flanery, Marci A.

Publication Date

2001

Peer reviewed

Prototype Abstraction in Category Learning?

Thomas J. Palmeri (thomas.j.palmeri@vanderbilt.edu)
Marcia A. Flanery (marcia.flanery@vanderbilt.edu)
Department of Psychology; Vanderbilt University
Nashville, TN 37240 USA

Abstract

Do people learn categories by abstracting prototypes, forming simple rules, remembering specific exemplars, or by some combination of these? Although some consensus seems to be emerging for a combination of rule formation and exemplar memorization, recent work has revived interest in prototype abstraction (e.g., Smith et al., 1997; Smith & Minda, 1998). We reexamined this recent evidence with an eye toward an alternative simple strategy subjects could use within those particular studies. A very simple strategy, available in some categorization tasks in which connective feedback is supplied, is to classify the current stimulus in the same category as the previous stimulus if the two are sufficiently similar to one another. This simple strategy makes no recourse to stored category representations of any kind. And this strategy will be useful only under certain circumstances. Reexamining the work by Smith and colleagues, we found that those category structures that produced evidence for prototype abstraction could be "learned," at least to some degree, using this simple strategy. Moreover, simulated data sets created using this simple strategy were better fitted by a prototype model than an exemplar model. We argue that evidence for prototype abstraction from the studies by Smith and colleagues may be weaker than they originally claimed.

Introduction

Do people learn categories by abstracting prototypes, forming rules, remembering exemplars, or some combination of these? In the domain of learning novel perceptual categories, we can trace an evolving dominance of various theoretical accounts from rule formation via hypothesis testing in the early years of categorization research (e.g., Bruner et al., 1956; Trabasso & Bower, 1968), to prototype abstraction (e.g., Homa, 1984; Posner & Keele, 1968; Reed, 1972), to exemplar storage and retrieval (e.g., Hintzman, 1986; Medin & Schaffer, 1978; Nosofsky, 1986). More recently, there has been a reemerging interest in the importance of rule learning in categorization (e.g., Allen & Brooks, 1991; Nosofsky et al., 1994; Nosofsky & Palmeri, 1998). This has led to a variety of hybrid accounts proposing a combination of rule learning and exemplar memorization (e.g., Erickson & Kruschke, 1998; Johansen & Palmeri, 2001; Palmeri, 1997; Palmeri & Johansen, 1999; see also Ashby et al., 1998). Arguably, there seems to be an emerging consensus for some kind of combination of

rule formation and exemplar memorization in category learning. However, there has also been reemerging interest in the potential role of prototype abstraction in category learning as well, at least under certain conditions (e.g., Smith et al., 1997; Smith & Minda, 1998). The goal of the present article was to critically reexamine this evidence for prototype abstraction that has been provided by Smith and colleagues.

Evidence for Prototype Abstraction

According to prototype models, people learn categories by averaging their experiences with specific exemplars to derive an abstract prototype and classify new objects according to similarity to stored prototypes. By contrast, according to exemplar models, people remember information about specific exemplars, with no abstraction of prototypes or other summary representations, and classify new objects according to their similarity to the stored category exemplars. Numerous studies have compared and contrasted the ability of exemplar models and prototype models to account for observed categorization data – the majority of studies found that exemplar models provided a far superior account, both qualitatively and quantitatively (e.g., Buse-

Table 1: An example category structure from Smith and Minda (1998), Experiments 1 and 2.

Category A		Category B	
Structure	Stimuli	Structure	Stimuli
Categories Linearly Separable			
000000	banuly	111111	kepio
010000	benuly	111101	kepib
100000	kanuly	110111	kenio
000101	banib	101110	kapiy
100001	kanub	011110	bepiy
001010	bapury	101011	kapuro
011000	bepuly	010111	benio
Categories Not Linearly Separable			
000000	gafuzi	111111	wysero
100000	wafuzi	011111	gysero
010000	gyfuzi	101111	wasero
001000	gasuzi	110111	wyfero
000010	gafuri	111011	wysuro
000001	gafuzo	111110	wyseri
111101	wysezo	000100	gafezi

meyer et al., 1984; Medin & Schaffer, 1978; Palmieri & Nosofsky, 2001; Shin & Nosofsky, 1992).

Challenging this previous work, there have been some recent studies that have reexamined the potential explanatory power of prototype models compared to exemplar models under a variety of conditions. In this work, prototype models have been found to provide superior fits to some data (Smith et al., 1997), especially early in category learning (Smith & Minda, 1998).

Smith et al. (1997) conducted a series of category learning experiments using structures and stimuli quite similar to those shown in Table 1. Stimuli were six-letter (6D) pronounceable nonsense words composed of alternating consonants and vowels. At each position in the word, one of two possible letters could appear (e.g., the first letter could be either g or w). The two categories were generally formed around a prototype (e.g., gafuzi versus wysezo) with most category members differing from their category prototype along one or two dimensions (e.g., gasuzi). Category structures were either linearly separable (LS) or nonlinearly separable (NLS). Linearly separable categories are those that can be partitioned on the basis of some weighted additive combination of information along the individual dimensions (Medin & Schwanenflugel, 1981). As shown in Table 1, the nonlinearly separable categories have an exception item that is similar to the prototype of the contrasting category (e.g., wysezo in the gafuzi category). In the first experiment of Smith et al. (1997), one set of category structures was relatively easy (ELS and ENLS) and another set of category structures was relatively difficult (DLS and DNLS).

Each subject learned one of the four possible category structures. On each trial of the experiment, the subject was presented with one of the items randomly selected from one of the two categories to be learned. The subject classified the item as a member of category A or category B and then received corrective feedback.

Smith et al. (1997) tested the ability of prototype and exemplar models to account for the probabilities of classifying each stimulus as a member of category A or B on an individual subject basis. A cross all four conditions of Experiment 1, they found that the prototype model provided a better account of the observed data than the exemplar model for half of their subjects. Their evidence for a prototype model advantage is summarized in Table 2. A YES in the second column signifies that at least some proportion of the individual subjects were displaying categorization behavior that a prototype model was better able to account for.

To illustrate, focusing on the ENLS category, the subgroup of subjects that a prototype model better accounted for averaged 92%, 78%, and 23% correct on the prototypes, normal items, and exception items, respectively. Consistent with the predictions of the prototype model, these subjects erroneously classified the exception items as being members of the category of the prototype they were most similar to. By contrast, the

subgroup of subjects that an exemplar model better accounted for averaged 81%, 79%, and 51% correct on the prototypes, normal items, and exceptions.

Smith et al. (1997) and Smith and Minda (1998) provided further evidence for a prototype model advantage across a series of experiments. One manipulation they performed varied the number of dimensions present in the stimuli. With six dimensions (6D) it is possible to create well-differentiated categories (those with high within-category similarity and low between-category similarity) with many members. However, with only four stimulus dimensions (4D), categories tend to be much less differentiated and tend to be much smaller. Smith and colleagues have argued that prototype models show their advantage where prototype abstraction is most likely to succeed, under those conditions where categories are composed of stimuli with many dimensions, where categories are large in size, and where categories are well differentiated. (As a further manipulation, in some experiments nonsense words were used, while in other experiments cartoon animals were used.) As summarized in Table 2, Smith et al. (1997) documented a series of conditions under which some proportion of subjects used prototype abstraction. In Smith and Minda (1998), evidence for prototype abstraction, if present, was observed in the early stages of category learning; exemplar models generally fared better than prototype models in later stages of learning. For experiments from that article, a YES in the second column of Table 2 signifies that a prototype model provided a superior account of early stages of category learning. As shown in Table 2, for some of the category structures Smith and colleagues tested, no evidence for prototype abstraction was observed.

A Simple Categorization Strategy

In all of the experiments cited in Table 2, prototype and exemplar models were tested on their ability to account for category learning data. This data was obtained from trials in which subjects were presented a stimulus, made a response, received feedback, were presented the next stimulus, made a response, received feedback, and so on. Our goal was to investigate whether some subjects could be using some form of the following very simple strategy to provide the correct answer without relying on abstracted prototypes or learned exemplars.

We will use the category structure shown in Table 1 as an example. Suppose on some trial, a subject is shown the following stimulus

gafuzi

and is then asked to classify it as a member of category A or category B. The subject responds *category A* and the computer provides the following feedback

CORRECT, gafuzi is a member of *category A*
Suppose the subject is next presented this stimulus

gafuri

and is asked to classify it. The subject could rely on abstracted prototypes, or remembered exemplars, or

formed rules. But perhaps a far simpler strategy is to classify it in the same category as *gafuzi* since they are so similar to one another. The computer just verified that *gafuzi* is a member of *category A* so it is reasonable to guess that *gafuzim* might also be a member of *category A* as well. The subject responds *category A* and the computer provides the following feedback

CORRECT, *gafuzim* is a member of *category A*
 Suppose the subject is next presented this stimulus

waseio
 and is asked to classify it. Well, this stimulus looks quite different from the previous stimulus, *gafuzi*, so it might make sense to classify it in the opposite category as *gafuzi*. The subject responds *category B* and the computer provides the following feedback

CORRECT, *waseio* is a member of *category B*
 Finally, suppose the subject is presented this stimulus
gafezi

Well, this stimulus looks very different from the previous one, *waseio*, so it might make sense to classify it in the opposite category. The subject responds *category A* and the computer provides the following feedback

WRONG, *gafezi* is a member of *category B*
 Examining Table 1, we see that this stimulus is the exception to category B. Using this very simple strategy, our subject would seem to perform quite well at classifying everything but the exceptions. Recall that Smith et al. (1997) reported that their subjects whose data was best fit by a prototype model consistently classified the exceptions in the wrong category as well.

Table 2 : Evidence for Prototype Abstraction from Smith et al. (1997) and Smith & Minda (1998). See text for a key to the experiment notation.

Experiment	Prototypes?	Strategy?
Smith et al. (1997)		
Experiment 1 6D ELS	YES	YES
Experiment 1 6D DLS	YES	YES
Experiment 1 6D ENLS	YES	YES
Experiment 1 6D DNLS	YES	YES
Experiment 2 4D NLS	NO	NO
Experiment 2 6D NLS	YES	YES
Smith & Minda (1998)		
Experiment 1 6D LS	YES	YES
Experiment 1 6D NLS	YES	YES
Experiment 2 6D LS	YES	YES
Experiment 2 6D NLS	YES	YES
Experiment 3 4D LS	NO	NO
Experiment 3 4D NLS	NO	NO
Experiment 4 4D NLS	NO	NO
Experiment 4 6D NLS	YES	YES

Note. The second column (Prototypes) indicates whether evidence for prototype abstraction was observed. The third column (Strategy) indicates whether the simple strategy yields above chance performance.

Does the Simple Strategy Work?

This is certainly a strategy that subjects could use to classify stimuli in an experiment. But, does it really work? In many cases, no. For example, let us consider the experiments reported by Medin and Schwanenflugel (1981). Different groups of subjects learned linearly separable and nonlinearly separable categories. Their results were important because they found that NLS categories could be easier to learn than LS categories, a result inconsistent with additive prototype models. By contrast, this result was an a priori prediction of multiplicative exemplar models. Let us first examine the category structure from the third experiment from Medin and Schwanenflugel (1981) in some detail. Using an abstract notation, for the LS structure, stimuli in category A were 0111, 1110, and 1001 and stimuli in category B were 1000, 0001, and 0110. For the NLS structure, stimuli in category A were 1100, 0011, and 1111, and stimuli in category B were 0000, 0101, 1010. We performed a Monte Carlo simulation of the simple strategy using these two category structures. For each of 1000 simulated subjects for each structure, we generated a random sequence of stimulus trials. On each simulated trial, if the current stimulus matched the previous one on more than two dimensions, then the same category response as the previous stimulus was used. If the current stimulus matched the previous one on fewer than two dimensions, then the other category response was used. If the current stimulus matched the previous one on exactly two dimensions, then a random response

Table 3 : Best fitting model (exemplar or prototype) to data simulated using the simple strategy. See text for a key to the experiment notation.

Experiment	Prototypes?	Model
Smith et al. (1997)		
Experiment 1 6D ELS	YES	Prototype
Experiment 1 6D DLS	YES	Prototype
Experiment 1 6D ENLS	YES	Prototype
Experiment 1 6D DNLS	YES	Prototype
Experiment 2 4D NLS	NO	—
Experiment 2 6D NLS	YES	Prototype
Smith & Minda (1998)		
Experiment 1 6D LS	YES	Prototype
Experiment 1 6D NLS	YES	Prototype
Experiment 2 6D LS	YES	Prototype
Experiment 2 6D NLS	YES	Prototype
Experiment 3 4D LS	NO	—
Experiment 3 4D NLS	NO	—
Experiment 4 4D NLS	NO	—
Experiment 4 6D NLS	YES	Prototype

Note. The second column (Prototypes) indicates whether evidence for prototype abstraction was observed. The third column (Model) indicates whether the Prototype or Exemplar model provided a better fit to the simulated data.

was generated. Averaging across 1000 simulated subjects, this strategy produced just 34.1% accuracy on the LS structure and 33.7% accuracy on the NLS structure. To see why this simple strategy produced accuracy far worse than just guessing, let us examine the NLS structure. Both NLS categories contain stimuli that mismatch each other on every dimension (1100 and 0011 in category A, 0101 and 1010 in category B). When these mismatching stimuli follow one another, they always produce the wrong response (e.g., erroneously categorizing 0011 as a member of category B because it is preceded by 1100 which was labeled a member of category A). Moreover, on other pairs of sequential trials, stimuli that follow one another match on exactly half the dimensions, producing a random response.

This simple strategy fails at other category structures as well. For the second experiment of Medin and Schwanenflugel (1981), the simple strategy produced 46.6% accuracy for their LS structure. For the category structure from Experiment 4 of Medin and Schaffer (1978), the simple strategy produced 44.6% accuracy. Applying the simple strategy to the classic category structures from Shepard, Hovland, and Jenkins (1961), we obtain predicted accuracies for their problem Types I-IV as following: Type I : 70.8% , Type II : 42.7% , Type III : 57.3% , Type IV : 57.0% , Type V : 43.1% , Type VI : 15.2% . In addition to underpredicting the overall level of accuracy observed when subjects learn these various categories, this simple strategy mispredicts the order of difficulty of the various problem types. For separable-dimension stimuli, the difficulty of the problems is ordered $I < II < III, IV, V < VI$ (Nosofsky et al., 1994; see also Nosofsky & Palmieri, 1996). Clearly, this simple strategy is not what subjects can use in many categorization experiments.

But, the simple strategy does work well when "learning" other category structures. Let us turn now to the category structures shown in Table 1, which were used in Experiments 1 and 2 of Smith and Minda (1998). Following the procedure described above, we used a Monte Carlo simulation procedure to generate categorization responses for the LS and NLS categories using the simple strategy. For each of 1000 simulated subjects, we generated a random sequence of stimuli, with each stimulus presented once per block. On each trial, if the current stimulus matched the previous one on more than three dimensions, then the category of the previous stimulus was used. If the current stimulus matched the previous one on fewer than three dimensions, then the other category was used. If the current stimulus matched the previous one on exactly three dimensions, then a random response was generated.

Using this simple strategy, accuracy of approximately 73% correct was possible for both the LS and NLS category structures (excluding the two exceptions in the NLS structure, which the strategy erroneously classified, the overall accuracy for the remaining items was over 84%). The overall performance is less than the ac-

curacies achieved by subjects in Smith and Minda (1998) by the end of learning, which was slightly over 80% correct for both structures. However, in their experiment, evidence for the use of prototypes was only observed during the early blocks of learning. Smith and Minda fitted the prototype and exemplar models to blocks of 56 trials and found that the prototype model fitted better than the exemplar model during the early blocks of learning. It seems quite possible that subjects might start out using the simple strategy during the early blocks of learning, especially since the strategy correctly classifies nearly three out of four items. As subjects acquire more experience with the categories, they may begin to shift to using stored exemplar information to improve performance.

We next tested the ability of this simple strategy to correctly categorize stimuli from the other category structures used by Smith and colleagues. As summarized in Table 2, the simple strategy produced above chance categorization in just those category structures that Smith and colleagues found evidence for prototype abstraction. For notation, those category structures for which the simple strategy works are indicated by a YES in the third column of Table 2. Could the apparent use of prototypes actually be a signature for the use of this very simple strategy instead?

Which Model Fits Better?

Suppose that subjects are engaging in this simple strategy of comparing the current stimulus with the previous one and selecting the category label accordingly. Let us further assume that they are not abstracting prototypes, are not learning rules, and are not remembering exemplars. Smith and colleagues obtained data from their subjects and compared how well a prototype model and an exemplar model accounted for categorization judgments. If subjects are using the simple strategy, would a prototype model or an exemplar model provide a better account of the categorization judgments produced using this simple strategy?

We will focus on just those structures for which the simple strategy actually yields above chance categorization, as indicated by a YES in the third column of Table 2. Using the simple strategy, we employed the Monte Carlo simulation techniques discussed earlier to generate data from a large number of simulated subjects for each of the indicated category structures. We then examined how well a prototype model or an exemplar model accounted for this simulated data. Clearly one possibility is that the prototype model accounted for some of the simulated data and an exemplar model accounted for the rest of the simulated data. A far more interesting possibility is that either an exemplar model or a prototype model provided a better account for the entire set of simulated data. This would pose an interesting problem of identifiability. The data were generated using a simple strategy of local stimulus comparisons without storing long-term category representations of

any kind. Yet by comparing just a prototype model and an exemplar model, we may erroneously conclude on the basis of model fits that subjects were actually abstracting prototypes or remembering exemplars.

To be specific, we fitted a prototype model and an exemplar model to the simulated data generated using the simple strategy. For the exemplar model, an item to be classified is compared with the stored exemplars of category A and category B. The probability of classifying the item into one of those categories is given by the relative summed similarity of the item to the stored exemplars of the two categories. For the prototype model, the probability of classifying the item is given by the relative similarity to the prototypes of the two categories. Similarity between an item i and a stored exemplar j (or an abstracted prototype j) is given by

$$S_{ij} = \prod_m s_m^{\delta(i,j)}$$

where the $0 < s_m < 1$ are free parameters along all m dimensions. The $\delta(i,j)$ is a function that returns a 0 if i and j match along dimension m and returns a 1 if i and j mismatch. A small value of s_m indicates that dimension m is particularly diagnostic. Because of the multiplicative similarity rule, mismatches along that dimension will have a large effect on decreasing similarity.

For the exemplar model, evidence for category A, E_A , is found by summing the similarities of an item to all exemplars in category A, and evidence for category B, E_B , is found by summing the similarities to all exemplars in category B. For the prototype model, evidence for category A, E_A , is the similarity to the category A prototype, and the evidence for category B, E_B , is the similarity to the category B prototype. The probability of classifying i into category A is then given by

$$P(A | i) = \frac{b_A E_A}{b_A E_A + b_B E_B}$$

where b_A is the category A bias (as might be expected, the bias terms did not contribute to any significantly improved fits to the simulated data). The prototype model and the exemplar model were fitted to the simulated data using a hill-climbing procedure that located parameters that minimized the sum of squared error between simulated observations and model predictions.

The summary of this modeling was straightforward. The prototype model provided a better account of the data that was simulated using the simple strategy than did the exemplar model. This finding is summarized in Table 3. As shown in the third column, for every structure which could be "learned" using the simple strategy, and for which Smith and colleagues found evidence for prototype abstraction, the prototype model provided a superior account of the simulated data. Although the data was generated using a strategy that made no recourse to abstract prototypes, the prototype model fitted that simulated data better than the exemplar model. If subjects were indeed using this simple strategy, one might erroneously conclude that they were abstracting

prototypes when in fact they were relying on local stimulus information to make a categorization decision.

Summary and Conclusions

Smith and colleagues (Smith et al., 1997; Smith & Minda, 1998) have provided evidence for prototype abstraction in category learning. We noted that in their experiments, they examined data from learning trials in which feedback was always provided. We investigated the possibility that subjects could use this corrective feedback to classify the subsequent item in the category learning task without making recourse to long-term category representations of any kind. According to this simple strategy, subjects must only compare the current item with the previous one. The previous item had been labeled explicitly by the experimenter in the form of corrective category feedback. The current item is classified in the same category as the previous one if they are sufficiently similar to one another, otherwise the current item is classified in the other category.

First, we observed that those structures that Smith and colleagues found evidence for prototype abstraction were those category structures which could be "learned" using this simple strategy. In other words, engaging in these local stimulus comparisons could produce categorization accuracy greater than chance. For comparison, we documented a number of other category structures for which this simple strategy would be unsuccessful.

Second, we simulated data using this simple strategy for those structures for which the simple strategy would work. We then fitted a prototype and an exemplar model to this simulated data. In every case, the prototype model fitted the simulated data better. If subjects were to use this simple strategy, without relying on stored category representations, we might erroneously conclude from these model fits that subjects were abstracting prototypes when they were not.

In a recent paper, Stewart, Brown, and Chater (in press) documented evidence for the use of sequential information in categorization similar to what we are proposing. Using what they called a memory and contact (MAC) strategy, they tested whether subjects would respond with the same category as on the previous trial if there was a small difference between the two stimuli, and a different label if the difference was large, just like the simple strategy we investigated. Not only did they demonstrate that the MAC strategy could achieve well over 80% accuracy on the category structures they tested, but they also reported highly systematic sequence effects in the experiments they reported, which were consistent with the use of a MAC strategy. The sequence effects they observed were inconsistent with exemplar models and other models they investigated.

So, the conclusion of our work along with that of Stewart and colleagues is perhaps best described as a cautionary tale. When we engage subjects in category learning experiments, our goal is typically to understand something about the long-term category representations

that subjects may have formed about those categories. Yet, subjects will use whatever information they have available to them to make a correct response, perhaps without even using any long-term category knowledge. They can clearly use feedback from previous trials to categorize on a current trial, at least under some circumstances. Or they can learn something about categories during testing in ways that may be entirely unanticipated by the investigator.

Acknowledgments

This work was supported by Vanderbilt University Research Council Grants, Grant BCS-9910756 from the National Science Foundation, and Grant 1R01MH61370 from the National Institute of Mental Health.

References

- Allen, S W., & Brooks, L R. (1991). Specializing the operation of an explicit rule. *Journal of Experimental Psychology: General*, *120*, 3-19.
- Ashby, F G., Alfonso-Reese, L A., Turken, A J., & Waldron, E M. (1998). A formal neuropsychological theory of multiple systems in category learning. *Psychological Review*, *105*, 442-481.
- Bruner, J S., Goodnow, J J., & Austin, G A. (1956). *A study of thinking*. New York: Wiley.
- Bussey, J R., Dewey, G I., & Medin, D L. (1984). Evaluation of exemplar-based generalization and the abstraction of categorical information. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *10*, 638-648.
- Erickson, M A., & Kruschke, J K. (1998). Rules and exemplars in category learning. *Journal of Experimental Psychology: General*, *127*, 107-140.
- Hintzman, D L. (1986). "Schemata abstraction" in a multiple-trace memory model. *Psychological Review*, *93*, 411-428.
- Homa, D. (1984). On the nature of categories. *Psychology of Learning and Motivation*, *18*, 49-94.
- Johansen, M K., & Palmeri, T J. (2001). Representational shifts in category learning. *Manuscript under review*.
- Medin, D L., & Schaffer, M M. (1978). A context theory of classification learning. *Psychological Review*, *85*, 207-238.
- Medin, D L., & Schwanenflugel, P J. (1981). Linear separability in classification learning. *Journal of Experimental Psychology: Human Learning and Memory*, *75*, 355-368.
- Nosofsky, R M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, *115*, 39-57.
- Nosofsky, R M., Gluck, M., Palmeri, T J., McKinley, S C., & Gauthier, P. (1994). Comparing models of rule-based classification learning: A replication and extension of Shepard, Hovland, and Jenkins (1961). *Memory & Cognition*, *22*, 352-369.
- Nosofsky, R M., & Palmeri, T J. (1996). Learning to classify integral-dimension stimuli. *Psychonomic Bulletin & Review*, *3*, 222-226.
- Nosofsky, R M., Palmeri, T J., & McKinley, S C. (1994). Rule-plus-exception model of classification learning. *Psychological Review*, *101*, 53-79.
- Nosofsky, R M., & Palmeri, T J. (1997). An exemplar-based random walk model of speeded classification. *Psychological Review*, *104*, 266-300.
- Nosofsky, R M., & Palmeri, T J. (1998). A rule-plus-exception model for classifying objects in continuous-dimension spaces. *Psychonomic Bulletin & Review*, *5*, 345-369.
- Palmeri, T J. (1997). Exemplar similarity and the development of automaticity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *23*, 324-354.
- Palmeri, T J., & Johansen, M K. (1999). Prototypes, rules, and instances in category learning. *Abstracts of the Psychonomic Society: 40th Annual Meeting*, *4*, 98.
- Palmeri, T J., & Nosofsky, R M. (2001). Central tendencies, extreme points, and prototype enhancement effects in ill-defined perceptual categorization. *The Quarterly Journal of Experimental Psychology*, *54*, 197-235.
- Posner, M I., & Keele, S W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, *77*, 353-363.
- Reed, S K. (1972). Pattern recognition and categorization. *Cognitive Psychology*, *3*, 382-407.
- Shepard, R N., Hovland, C L., & Jenkins, H M. (1961). *Psychological Monographs*, *75* (13, Whole No. 517).
- 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., Murray, M J., & Medin, J P. (1997). Straight talk about linear separability. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *23*, 659-680.
- Smith, J D., & Medin, J P. (1998). Prototypes in the mist: The early epochs of category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24*, 1411-1436.
- Stewart, N., Brown, G D A., & Chater, N. (in press). Sequence effects in categorization of simple perceptual stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition*.
- Trabasso, T., & Bower, G H. (1968). *Attention in learning: Theory and research*. New York: Wiley.