

The Time-Course of the Use of Background Knowledge in Perceptual Categorization

Thomas J. Palmeri (thomas.j.palmeri@vanderbilt.edu)

Department of Psychology; Vanderbilt University
Nashville, TN 37240 USA

Celina Blalock (ladybug531@juno.com)

Department of Child and Family Development; University of Georgia
Athens, GA 30604 USA

Abstract

We examined the time-course of the utilization of background knowledge in perceptual categorization by manipulating the meaningfulness of labels associated with categories and by manipulating the amount of time given to subjects to make a categorization decision. Extending a paradigm originally reported by Wisniewski and Medin (1994), subjects learned two categories of children's drawings that either were given standard labels (drawing by children from group 1 or group 2) or were given theory-based labels (drawings by creative or noncreative children); meaningfulness of the label had a profound effect on how new drawings were categorized. Half of the subjects were given unlimited time to respond, the other half of the subjects were given a quick response deadline; speeded response conditions had a relatively large effect on categorization decisions by subjects given the standard labels but had a relatively small effect on categorization decisions by subjects given the theory-based labels. These results suggest that background knowledge may have its influence at relatively early stages in the time-course of a categorization decision.

Introduction

In cognitive science, two distinct approaches to research on categorization and concept formation have been considered. One approach focuses on the statistical aspects of learning categories via induction over a series of instances. A main theme of this work involves developing and testing formal mathematical models of categorization that embody various assumptions about the kinds of representations thought to be formed during category learning. Examples include rule-based models (e.g., Nosofsky & Palmeri, 1998; Nosofsky, Palmeri, & McKinley, 1994; Palmeri & Nosofsky, 1995), exemplar-based models (Estes, 1994; Hintzman, 1986; Medin & Schaffer, 1978; Nosofsky, 1986; Nosofsky & Palmeri, 1997), various clustering models (e.g., Anderson, 1990; Fisher, 1987), and numerous others. What these theoretical approaches have in common is that the category representations that are formed during learning are entirely dependent on the empirical regularities in the particular set of category instances that have been experienced.

A contrasting approach focuses on how background knowledge or theories might influence what is learned about a category from experience with particular instances (Murphy & Medin, 1985). For example, background knowledge can influence the ease of learning linearly separable versus

nonlinearly separable categories (Wattenmaker, Dewey, Murphy, & Medin, 1986) and can influence the ease of learning conjunctive versus disjunctive rules (Pazzani, 1991). In addition, a number of studies have found a facilitative effect of prior background knowledge on learning new categories (e.g., Heit, 1994, 1998; Murphy & Allopenna, 1994; Murphy & Wisniewski, 1989).

We will describe the results of one particularly illuminating study by Wisniewski and Medin (1994), both to illustrate how background knowledge can influence categorization and because the present experiment builds on this particular study. Subjects learned two categories of drawings, shown in Figure 1. One group of subjects was provided standard labels (i.e., children from "group 1" drew the pictures in category A, children from "group 2" drew the pictures in category B), while another group of subjects was provided theory-based labels (i.e., "creative" children drew the pictures in category A, "noncreative" children drew the pictures in category B). Subjects were required to develop a set of rules that could be used to partition those drawings as well as any new drawings into the two distinct categories. When Wisniewski and Medin analyzed the rules subjects formed, they found that those subjects given the standard labels had generated rules based on fairly concrete perceptual features (such as "... all of the characters have their arms out straight from their bodies and they're also standing very straight, facing the front") whereas those subjects given the theory-based labels had generated rules based on fairly abstract features (such as "much more attention was given to the clothing ..." or "make drawings that show more positive emotional expression ...").

Wisniewski and Medin next tested their subjects on new drawings that factorially combined the perceptual features and the abstract features of the two categories, as shown in Figure 2. For example, stimulus T5 had the abstract features of category A but had the concrete perceptual features of category B, whereas stimulus T15 had the abstract features of category B but had the concrete perceptual features of category A. As shown in Figure 3, Wisniewski and Medin observed that subjects given different labels classified these conflicting stimuli in markedly different ways; subjects provided the theory-based labels tended to classify according to the abstract features and subjects provided the

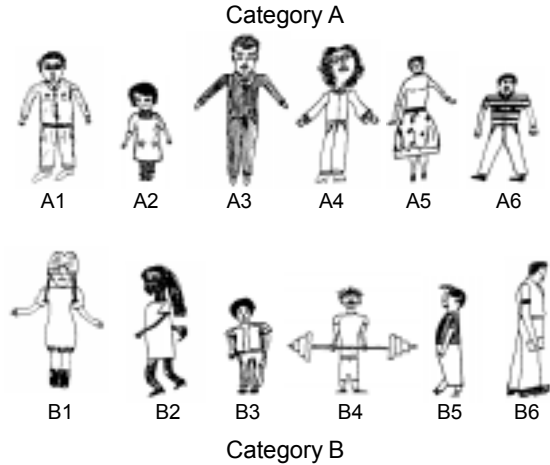


Figure 1. Training drawings used in the experiment (from Wisniewski & Medin, 1994).




Figure 2. Transfer drawings used in the experiment (from Wisniewski & Medin, 1994).




Figure 3. Probability of categorizing each of the transfer drawings into category A as a function of the type of label provided to subjects (data from Experiment 1A of Wisniewski and Medin, 1994).

standard labels tended to classify according to the concrete perceptual features. For example, note in Figure 3 how stimuli T5 and T15 were classified into different categories depending on what kind of label subjects were provided.

These results may be difficult to explain by theories of categorization that assume nothing more than some form of an abstraction of the statistical structure of the training patterns. In this experiment, both groups of subjects observed the identical set of training drawings. Yet, by simply manipulating the meaningfulness of the labels applied to those drawings, subjects categorized new transfer drawings in markedly different ways. To explain these results, Wisniewski and Medin (1994) suggested that background knowledge and empirical information about instances closely interact during categorization. They contended that background knowledge does not just weight the features of an object that are extracted during early perceptual processing. Rather, background knowledge may influence what features are actually extracted from an object for purposes of categorization (see also Wisniewski & Medin, 1991). The suggestion is that knowledge may have an influence on categorization at the very earliest stages of processing, not just after the objects have been analyzed into their constituent parts.

We chose to investigate this claim by manipulating the amount of time provided to subjects for making a categorization decision. We chose to introduce a response deadline in which one group of subjects was signaled to respond very soon after stimulus onset while another group of subjects was given unlimited time to respond. Such response signal paradigms can be quite useful for determining what information is available at various points within the time-course of a categorization decision (e.g., Lamberts, 1998; see also Meyer, Irwin, Osman, & Kounios, 1988).

For example, it is possible that background knowledge might influence categorization at relatively late stages of processing, contrary to the claims made by Wisniewski and

Medin (1994) – knowledge might be used to interpret features after they have been extracted, or knowledge might be used to guide the combination of a number of simple perceptual features into something more meaningful and useful for determining category membership. Introducing a response deadline might force subjects to make their categorization response before background knowledge has yet had any time to perform these interpretive operations. Therefore, under a response deadline, subjects given the theory-based labels might have much more difficulty categorizing both training and transfer drawings than subjects given the standard labels.

In this experiment, we extended the paradigm developed by Wisniewski and Medin (1994). Subjects learned the categories of drawings shown in Figure 1 either using standard labels or using theory-based labels. Subjects were then asked to classify the training drawings and new transfer drawings either under speeded conditions (issuing a response signal just 200ms after stimulus onset) or under unspeeded conditions (unlimited decision time).

Method

Subjects. Sixty Vanderbilt University undergraduates participated for partial course credit.

Stimuli. The stimuli were drawings by children who were given a "draw-a-person" test and were the same as those used in Experiment 1 of Wisniewski and Medin (1994; taken from Koppitz, 1984, and Harris, 1963). The drawings are displayed in Figures 1 and 2. For the first part of the experiment, the drawings were individually mounted and covered in clear-coat plastic. For the remaining parts of the experiment, scanned drawings were individually displayed on a computer monitor.

Following Wisniewski and Medin (1994), the training set consisted of two categories of six drawings each, as shown in Figure 1. The first set (category A) were deemed to be relatively detailed and the second set (category B) were deemed to be relatively unusual. In pilot studies Wisniewski and Medin conducted, people expected drawings made by creative children to be more detailed or more unusual than drawings made by noncreative children.

In addition to these abstract properties, the drawings in the two categories could be distinguished on the basis of at least three simple rules based on concrete perceptual features (Wisniewski and Medin modified the drawings to conform to these rules). For category A (top row of Figure 1), the drawings had "curly hair and arms not at the sides," "light-colored shoes or not smiling," and "wearing a collar or tie." For category B (bottom row of Figure 1), the drawings had "straight hair or arms at the sides," "dark-colored shoes and not smiling," or "ears and short sleeves."

Sixteen test drawings, shown in Figure 2, were selected and appropriately modified to factorially combine the abstract and concrete features of the training drawings. The first set of four drawings, shown in the top row of Figure 2, were not detailed and not unusual; the second set of four drawings, shown in the second row of Figure 2, were both detailed and unusual; the third set of four drawings, shown in the third row of Figure 2, were detailed and not unusual;

the fourth set of four drawings, shown in the bottom row of Figure 2, were not detailed and unusual. For each set of four drawings, two were best described by the concrete rules of the category A drawings and the other two were best described by the concrete rules of the category B drawings. For example, the first two drawings of the first set, test drawings T1 and T2, had the concrete rules of the category B, whereas the second two drawings, test drawings T3 and T4, had the concrete rules of category A. See Figure 2 for more information.

Procedure. Half of the subjects were instructed that they would be studying drawings by "creative" or "noncreative" children, the other half of the subjects were instructed that they would be studying drawings by children from "group 1" or "group 2." In the initial study phase, subjects were shown the mounted drawings simultaneously, laid out on the table in front of them, and separated into two labeled groups (with theory-based labels of "creative" or "noncreative", or standard labels of "group 1" or "group 2"). The instructions for the "creative" / "noncreative" subjects were as follows (adapted from Wisniewski and Medin, 1994):

In this experiment, you will be shown two groups of children's drawings. One group was done by creative children. The other group was done by noncreative children. Your task is to come up with a rule that someone could use to decide whether a new drawing belongs to the group drawn by the creative children. In writing down a rule, it is important that the rule "works" for all of the drawings in the creative group and none of the drawings in the noncreative group. That is, if someone were given your rule and the drawings from the two groups (all mixed up), they should be able to use it to divide the drawings into those that belong in one group and those that belong in the other group. Also, please write your rule clearly and describe the rule so that the experimenter will be able to understand what your mean.

The instructions for the "group 1" / "group 2" subjects were virtually the same, except that "creative" was replaced by "group 1" and "noncreative" was replaced by "group 2." Whereas Wisniewski and Medin (1994) randomized the assignment of category labels to sets of drawings, we instead chose to have the "creative" and "group 1" labels apply to the category A drawings and to have the "noncreative" and "group 2" labels apply to category B drawings.

After developing a set of rules to classify the drawings into the two categories, the mounted drawings were removed but subjects were allowed to keep their rules at hand. Subjects were then shown each of the training drawings one at a time on the computer, were asked to classify them as "creative" / "noncreative" or "group 1" / "group 2", were supplied corrective feedback, and were permitted to modify their rules if necessary. After classifying each of the drawings twice, their rules were removed. Order of drawings was randomized for every subject.

In the next phase of the experiment, subjects were asked to classify the training drawings and some new test drawings using the rules they had developed in the first part of the experiment. They were also instructed that they would

receive no corrective feedback following their response. Half of the subjects in each group were allowed to classify the drawings without any time limit (unspeeded condition), the other half of the subjects in each group were instructed that they would be required to make their responses according to a response signal (speeded condition).

In the speeded condition, a tone was presented 200ms after the onset of the drawing. Subjects were required to make their classification response within 300ms of hearing the tone. These subjects were provided a series of practice trials to familiarize themselves with the demands of the speeded condition. On the practice trials, a category label (either "creative" or "noncreative", or "group 1" or "group 2") was displayed at the center of the screen. After 200ms, the tone sounded and subjects were required to press the response key associated with the displayed category label. If the subject responded before the tone sounded, they were informed to wait until the tone sounded; if the subject responded more than 300ms after the tone sounded, they were informed to respond more quickly; if the subject made an incorrect response, they were informed not to make errors. The practice trials continued until the subject was able to make 15 valid responses in a row. Once the subject had achieved this criterion for responding appropriately to the signal, they were moved on to the test phase.

In the test phase, a crosshairs appeared at the center of the screen for 1 sec. Then one of the 28 drawings (12 training and 16 test) was displayed at the location of the crosshairs and the subject was required to classify it as "creative" / "noncreative" or "group 1" / "group 2". In the unspeeded condition, the subject could take as much time as they needed to make their response. In the speeded condition, a tone sounded 200ms after the drawing was displayed and subjects were required to make their response within 300ms of the tone; subjects were informed if they had made a response prior to the tone or had made a response more than 300ms after the tone. In both the speeded and unspeeded conditions, no other corrective feedback was provided. After the response was made, the drawing was erased from the screen. After a one second interval, the next drawing was presented. Subjects classified each of the 28 drawings eight times. Each stimulus was presented once per block in a randomized order for every subject.

To summarize the design of the experiment, each subject could be classified into one of four groups: theory-based labels (creative versus noncreative) with unspeeded responses, theory-based labels with speeded responses, standard labels (group 1 versus group 2) with unspeeded responses, and standard labels with speeded responses.

Results and Discussion

The primary data of interest were the probabilities of categorizing each stimulus as a member of category A in the testing phase as a function of category label (standard versus theory-based) and as a function of response deadline (speeded versus unspeeded). Figure 4 displays the probability of categorizing each of the original training drawings into category A. Figures 5 and 6 both display the probability of categorizing each of the new transfer drawings into

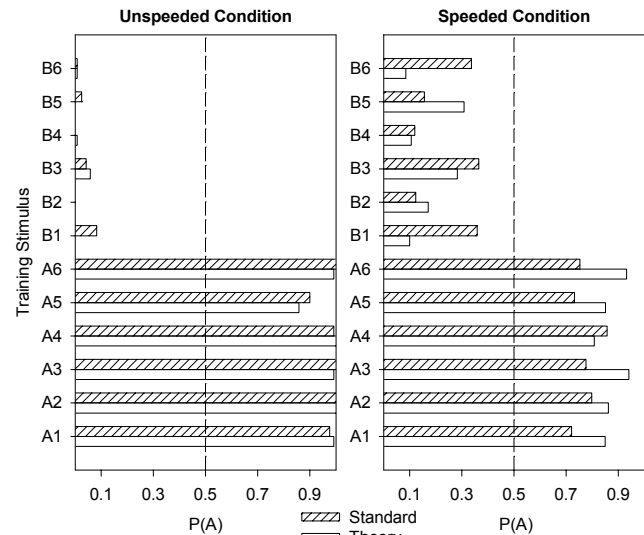


Figure 4. Probability of categorizing each of the training drawings into category A as a function of response deadline (unspeeded in the left panel, speeded in the right panel) and as a function of label (hatched bars for standard labels, open bars for theory-based labels).

category A (Figure 5 plots the two response deadline conditions separately, Figure 6 plots the two category label conditions separately). In the speeded condition, categorization responses that were made prior to the response signal or more than 300ms after the response signal were excluded from the analyses (12.8% of the observations); further analyses in which we excluded none of the categorization responses were essentially identical to those presented here.

A 28 (stimulus) \times 2 (category label) \times 2 (response deadline) mixed analysis of variance was conducted on the probability data with stimulus (12 training drawings and 16 transfer drawings) as a within-subjects variable, and with label (standard versus theory-based) and deadline (speeded versus unspeeded) as between-subjects variables. A significant main effects of stimulus, $F(27,1512)=56.454$, and significant interactions of stimulus \times label, stimulus \times deadline, and stimulus \times label \times deadline were observed, $F(27, 512)=18.282$, $F(27,1512)=5.787$, and $F(27,1512)=1.636$, respectively (an alpha level of .05 was established for all statistical tests reported in this article). No main effects of label or deadline were observed.

Not surprisingly, for the training drawings, subjects were significantly more accurate in the unspeeded condition than the speeded condition. Somewhat surprisingly, for the training drawings, subjects were significantly more accurate in the speeded condition when applying theory-based labels (84.1% correct) than when applying standard labels (75.6% correct). This provides our first piece of evidence inconsistent with the hypothesis that background knowledge might influence categorization only at relatively late stages in the decision making process.

Turning to the categorization of the transfer drawings, let us first note the fairly close correspondence between the findings we obtained in the unspeeded condition, shown in

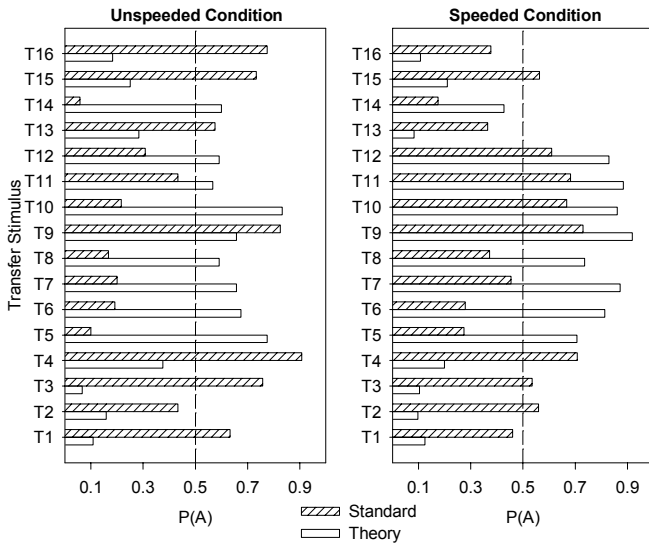


Figure 5. Probability of categorizing each of the transfer drawings into category A as a function of response deadline (unspeeded in the left panel, speeded in the right panel) and as a function of label (hatched bars for standard labels, open bars for theory-based labels).

the left panel of Figure 5, and the findings obtained by Wisniewski and Medin, shown in Figure 3. Our results provide an important replication of their study by again demonstrating that the meaningfulness of a label can have a striking effect on how subjects categorize new transfer drawings. The general pattern of how the transfer drawings were categorized can probably best be explained by assuming that subjects given theory-based labels tended to categorize on the basis of abstract features whereas subjects given standard labels tended to categorize on the basis of concrete perceptual features (see Wisniewski & Medin, 1994).

The novel aspect of the present work is a finding that response deadlines had quite different effects depending on the type of label that was provided. To see this, compare categorization probabilities under speeded and unspeeded conditions for subjects given standard labels, shown in the left panel of Figure 6, with those for subjects given theory-based labels, shown in the right panel of Figure 6. One useful way of summarizing this data is as follows: With standard labels, categorization probabilities tended to be closer to 50% or even crossed over to the opposite category in the speeded condition compared to the unspeeded condition. By contrast, with theory-based labels, categorization probabilities actually tended to be further from 50% in the speeded condition compared to the unspeeded condition. A simple explanation for this difference solely in terms of the amount of time needed to apply the rules in the two conditions is unlikely; in the unspeeded conditions, there was no significant difference in response times for subjects given standard (1420ms) and theory-based labels (1601ms).

Another useful way of summarizing this data requires slightly recoding the categorization responses as follows. To develop a measure analogous to categorization accuracy for the new transfer drawings, for each stimulus, we simply define the modal categorization response in the unspeeded

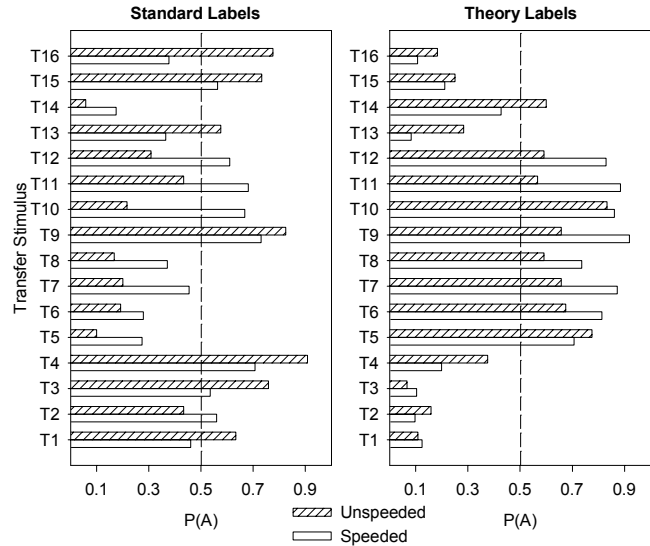


Figure 6. Probability of categorizing each of the transfer drawings into category A as a function of label (standard labels in the left panel, theory-based labels in the right panel) and as a function of response deadline (hatched bars for unspeeded, open bars for speeded).

condition as the "correct" category response for that stimulus (e.g., for subjects given standard labels, the "correct" response for T5 is category B, but for subjects given the theory-based labels, the "correct" response for T5 is category A). For each stimulus, we then recode each subjects' category responses with respect to this nominally "correct" response. Figure 7 displays categorization "accuracy" averaged across all the transfer stimuli. With standard labels, subjects were significantly less "accurate" in the speeded condition than the unspeeded condition; with theory-based labels, subjects actually tended to be more "accurate" in the speeded condition than the unspeeded condition, although this difference did not reach significance.

Either way we examine the data, we find that introducing a very rapid response deadline, causing subjects to make fairly complex categorization responses in an average of just 370ms, had surprisingly little effect on categorization performance when subjects were provided meaningful labels that tapped into background knowledge. Subjects

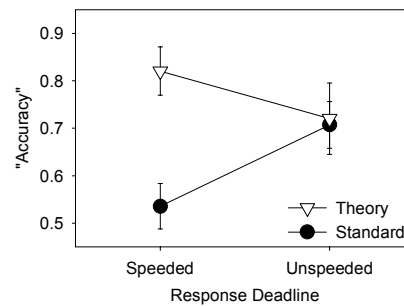


Figure 7. Categorization "accuracy" averaged across the transfer drawing as a function of response deadline (x-axis) and as a function of label (closed circles for standard labels, open triangles for theory-based labels).

appeared to be able to utilize this background knowledge at the earliest stages of a categorization decision, consistent with the hypothesis that the joint influences of background knowledge and empirical evidence are tightly coupled in the service of making a categorization decision (see Wisniewski, 1995). This is one of the first fairly direct demonstrations that background knowledge can have a very early influence in the time course of a perceptual categorization decision (see also Lin & Murphy, 1997).

It was somewhat surprising that introducing a very rapid response deadline had a quite a detrimental effect on categorization performance when subjects were provided standard labels. Although the rules subjects developed in this condition were based on simple perceptual features, testing these rules on a given drawing may have required a serial evaluation of the various component features making up that categorization rule (see Palmeri & Nosofsky, 1995). Under a response deadline, subjects simply did not have sufficient time to check all of the various features of the rule they formed, causing them to respond in a haphazard fashion.

Acknowledgments

The research described in this paper is based on work completed in partial fulfillment of an undergraduate honors thesis submitted to Vanderbilt University by the second author. This work was supported by Vanderbilt University Research Council Direct Research Support grants to the first author. We thank Douglas Medin and Evan Heit for their comments on this work.

References

- Anderson, J.R. (1990). *The adaptive character of thought*. Hillsdale, NJ: Erlbaum.
- Estes, W.K. (1994). *Classification and cognition*. Oxford, England: Oxford University Press.
- Fisher, D.H. (1987). Knowledge acquisition via incremental conceptual clustering. *Machine Learning*, 2, 139-172.
- Gluck, M.A., & Bower, G.H. (1988). Evaluating an adaptive network model of human learning. *Journal of Memory and Language*, 27, 166-195.
- Harris, D.B. (1963). *Children's drawings as measures of intellectual maturity*. New York: Harcourt Brace & World.
- Heit, E. (1994). Models of the effects of prior knowledge on category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20, 1264-1282.
- Heit, E. (1998) Influences of prior knowledge on selective weighting of category members. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24, 712-731.
- Koppitz, E.M. (1984). *Psychological evaluation of human figure drawings by middle school pupils*. Orlando, FL: Grune & Stratton.
- Lamberts, K. (1998). The time course of categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24, 695-711.
- Lin, E.L., & Murphy, G.L. (1997) Effects of background knowledge on object categorization and part detection. *Journal of Experimental Psychology: Human Perception and Performance*, 23, 1153-1169.
- Medin, D.L., & Schaffer, M.M. (1978). A context theory of classification learning. *Psychological Review*, 85, 207-238.
- Meyer, D.E., Irwin, D.E., Osman, A.M., Kounios, J. (1988). The dynamics of cognition: Mental processes inferred from a speed-accuracy decomposition technique. *Psychological Review*, 95, 183-237.
- Murphy, G.L., & Allopenna, P.D. (1994). The locus of knowledge effects in concept learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 20, 904-919.
- Murphy, G.L., & Medin, D.L. (1985). The role of theories in conceptual coherence. *Psychological Review*, 92, 289-316.
- Murphy, G.L., & Wisniewski, E.J. (1989). Feature correlations in conceptual representations. In G. Tiberghien (Ed.), *Advances in cognitive science: Vol. 2. Theory and applications*. Chichester, England: Ellis Horwood.
- Nosofsky, R.M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General*, 115, 39-57.
- 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 classify objects in continuous-dimension spaces. *Psychonomic Bulletin & Review*, 5, 345-369.
- 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.
- Pazzani, M.J. (1991). Influence of prior knowledge on concept acquisition: Experimental and computational results. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15, 416-432.
- Wattenmaker, W.D., Dewey, G.I., Murphy, T.D., & Medin, D.L. (1986). Linear separability and concept learning: Context, relational properties, and concept naturalness. *Cognitive Psychology*, 18, 158-194.
- Wisniewski, E.J. (1995). Prior knowledge and functionally relevant features in concept learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21, 449-468.
- Wisniewski, E.J., & Medin, D.L. (1991). Harpoons and long sticks: The interaction of theory and similarity in rule induction. In D.H. Fisher, M.J. Pazzani, & P. Langley (Eds.), *Concept formation: Knowledge and experience in unsupervised learning*. San Mateo, CA: Morgan Kaufmann.
- Wisniewski, E.J., & Medin, D.L. (1994). On the interaction of theory and data in concept learning. *Cognitive Science*, 18, 221-281.