



Brief article

The role of background knowledge in speeded perceptual categorization

Thomas J. Palmeri^{a,*}, Celina Blalock^b^a*Department of Psychology, 301 Wilson Hall, Vanderbilt University, Nashville, TN 37240, USA*^b*University of Georgia, Athens, GA, USA*

Received 27 October 1999; received in revised form 24 May 2000; accepted 16 June 2000

Abstract

We examined the time-course of the influence of background knowledge on perceptual categorization by manipulating the meaningfulness of labels associated with categories and by manipulating the amount of time provided to subjects for making a categorization decision. Extending a paradigm originally reported by Wisniewski and Medin (1994) (*Cog. Sci.* 18 (1994) 221), subjects learned two categories of children's drawings that were given either meaningless labels (drawings by children from 'group 1' or 'group 2') or meaningful labels (drawings by 'creative' or 'non-creative' children); the meaningfulness of the label had a significant effect on how new drawings were classified. In addition, half of the subjects were provided unlimited time to respond, while the other half of the subjects were forced to respond quickly; speeded response conditions had a relatively large effect on categorization decisions by subjects given the meaningless labels but had relatively little effect on categorization decisions by subjects given the meaningful labels. These results suggest that some forms of background knowledge can show an influence at relatively early stages in the time-course of a categorization decision. © 2000 Elsevier Science B.V. All rights reserved.

Keywords: Background knowledge; Speeded perceptual categorization

1. Introduction

An important topic in the study of categorization and concept formation focuses on how background knowledge or theories might influence what is learned about a category (Murphy & Medin, 1985). For example, background knowledge can influence the ease of learning linearly separable versus non-linearly separable categories

* Corresponding author. Tel.: +1-615-343-7900; fax: +1-615-343-8449.

E-mail address: thomas.j.palmeri@vanderbilt.edu (T.J. Palmeri).

(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 begin by describing 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 earlier study (from this point forward, this Wisniewski and Medin study will be referred to by the abbreviation W&M). In their experiment, subjects learned two categories of drawings, shown in the top half of Fig. 1 (Training Drawings). One group of subjects was provided standard meaningless labels (i.e. children from ‘group 1’ drew pictures in category A, children from ‘group 2’ drew pictures in category B), while another group of subjects was provided meaningful labels that tapped into a particular kind of background knowledge (i.e. ‘creative’ children drew pictures in category A, ‘non-creative’ children drew pictures in category B). Subjects were required to develop a set of written rules to partition the training drawings as well as any possible new drawings into the two categories. When analyzing the rules subjects formed, W&M found that subjects given meaningless labels generated rules based on concrete perceptual features (e.g. “...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 meaningful labels generated rules based on fairly abstract properties (e.g. “much more attention was given to the clothing...” or “make drawings that show more positive emotional expression...”).

W&M next tested subjects on new drawings that systematically combined perceptual features and abstract properties of the two categories, as shown in the bottom half of Fig. 1 (Transfer Drawings). For example, T13 has the abstract properties of category B and the concrete perceptual features of category B, whereas T10 has the abstract properties of category A but the concrete perceptual features of category B. W&M observed that subjects classified transfer drawings that conflicted on abstract properties and perceptual features in a manner consistent with the type of labels they were provided. Subjects provided meaningful labels classified according to abstract properties (e.g. tending to classify T10 and T13 as members of different categories), whereas subjects provided meaningless labels classified according to perceptual features (e.g. tending to classify T10 and T13 as members of the same category).

These results may be difficult to explain by theories of categorization that just assume an abstraction of the statistical structure of training patterns via rules (e.g. Nosofsky & Palmeri, 1998; Nosofsky, Palmeri, & McKinley, 1994), prototypes (e.g. Homa, 1984; Reed, 1972), or specific instances (e.g. Hintzman, 1986; Medin & Schaffer, 1978; Nosofsky, 1986). Both groups of subjects observed an identical set of training drawings that were divided into the same two categories. Yet, by manipulating the meaningfulness of the labels applied to those categories of drawings, subjects classified new drawings in markedly different ways. To explain these results, W&M suggested that background knowledge and empirical information about instances closely interact during category learning. They proposed that back-

Training Drawings


Category A (Detailed)



Category B (Unusual)

Transfer Drawings

not Detailed & not Unusual



T1 T2
Concrete rules of
Category B



T3 T4
Concrete rules of
Category A

Detailed and not Unusual



T9 T10 T11 T12
Concrete rules of
Category B Concrete rules of
Category A

Detailed & Unusual



T5 T6
Concrete rules of
Category B



T7 T8
Concrete rules of
Category A

not Detailed and Unusual



T13 T14 T15 T16
Concrete rules of
Category B Concrete rules of
Category A

Fig. 1. Training and transfer drawings used in the experiment (adapted from Wisniewski & Medin, 1994). The top portion of the figure displays the 12 training drawings. Category A drawings have the abstract property of being relatively ‘detailed’ and could be classified using simple perceptual rules such as ‘curly hair and arms not at the sides’, or ‘light-colored shoes or smiling’, or ‘wearing a collar or tie’. Category B drawings have the abstract property of being relatively ‘unusual’ and could be classified using simple perceptual rules such as ‘straight hair or arms at the sides’, or ‘dark-colored shoes and not smiling’, or ‘ears and short sleeves’. The bottom portion of the figure displays the 16 transfer drawings. The first set (T1–T4) have abstract properties of neither category, the second set (T5–T8) have the abstract properties of both categories, the third set (T9–T12) have the abstract properties of category A, and the fourth set (T13–T16) have the abstract properties of category B. Within each set of transfer drawings, two drawings follow the simple perceptual rules of category A and two drawings follow the simple perceptual rules of category B.

ground knowledge may not just weight the features of an object that are extracted during early perceptual processing, but may actually influence the nature of the features that are extracted (or possibly created) for purposes of categorization.

W&M did not specifically address when background knowledge might have its influence within the time-course of a categorization decision. However, if they are correct that background knowledge might influence how objects are perceptually processed, then the effects of different types of category labels on categorization decisions might emerge at fairly early stages of processing, not just after the objects have been perceptually analyzed into their constituent parts.

Some recent evidence for an early locus of knowledge effects comes from a series of experiments by Lin and Murphy (1997) in which subjects classified objects that were consistent or inconsistent with background knowledge under speeded or unspeeded conditions. For example, in one experiment, subjects were provided a description of how some novel object was used by a novel cultural group. In one condition, the object might be described as a hunting implement, with the description leading subjects to infer some importance for those parts of the object used to catch an animal for food; in another condition, the same object might be described as a tool for spraying pesticides, with the description leading subjects to infer some importance for other parts of the object. Test objects were provided that contained the critical parts consistent with the description provided to one group but not the other group (e.g. the object might have the parts necessary for hunting but not the parts necessary for pesticide application). Lin and Murphy (1997) observed that even under speeded conditions, subjects classified new test objects according to whether they possessed the features that were critical in order to carry out the function provided in the initial description.

We chose to further investigate the possibility that background knowledge could have an early influence on categorization by manipulating the amount of time provided to subjects for making a categorization decision within the W&M paradigm. We introduced a response deadline in which one group of subjects was signaled to respond very soon after stimulus onset while another group of subjects was provided unlimited time to respond. Such response signal paradigms can be useful for ascertaining what information is available at various time points within a categorization decision (e.g. Lamberts, 1998). For example, it is quite possible that the use of background knowledge might involve some kind of reflective rule-based process that influences categorization only when concerns about making relatively rapid responses are absent (e.g. Smith & Sloman, 1994). Knowledge might not influence early perceptual processing directly. Rather, knowledge might be used to interpret features after they have been extracted or knowledge might be used to guide the combination of simple perceptual features into something more meaningful and useful for categorization. If that were the case, then introducing a response deadline might force subjects to generate a categorization response before they have had time to use background knowledge in such a reflective manner to appropriately interpret or combine the perceptual features. Under a response deadline, subjects given meaningful labels might have much more difficulty categorizing drawings than subjects given meaningless labels, or at least show little influence of

background knowledge influencing how drawings were classified. By contrast, if background knowledge is used to guide perceptual analysis, or at least has some important influence on relatively early stages of the categorization process, then subjects put under a response deadline who are provided meaningful labels might categorize drawings in a qualitatively similar way as subjects given unlimited time to respond.

In this experiment, we extended the paradigm developed by W&M. Subjects learned categories of drawings shown in the top half of Fig. 1 either using meaningless labels or using meaningful labels. Subjects were then asked to categorize all of the drawings shown in Fig. 1 either under speeded conditions (a response signal was presented 200 ms after stimulus onset) or under unspeeded conditions (unlimited decision time). If knowledge operates only as a slow, optional process during the categorization decision, then speeded decisions should show much less influence of background knowledge than unspeeded decisions. By contrast, if knowledge influences early stages of a categorization decision, then speeded decisions should show qualitatively similar influences of background knowledge as unspeeded decisions.

Our work builds on Lin and Murphy (1997) in some important ways. First, our subjects were asked to categorize drawings following a signal presented just 200 ms after stimulus onset (with a maximum response time of 500 ms), whereas Lin and Murphy (1997) provided up to 1 s to categorize their stimuli. Second, Lin and Murphy (1997) used novel objects that had distinct parts that were explicitly labeled and which could be provided an explicit function; new test objects either had those critical parts or they did not. By contrast, although the simple perceptual features used in our study had been explicitly manipulated to some extent (see W&M), the abstract properties were less identifiable. That is, there were generally no explicit parts of a drawing that could be used to decide whether a drawing was by a creative or non-creative child. Thus, our work provides an important extension of Lin and Murphy (1997) and of W&M by testing subjects under extreme conditions of speed stress and using stimuli with more abstract descriptions not easily amenable to decomposition into basic parts.

2. Method

2.1. Subjects

Sixty Vanderbilt University undergraduates participated for partial course credit.

2.2. Stimuli

Stimuli were drawings by children who were given a ‘draw-a-person’ test and were the same as those used in Experiment 1 of W&M (adapted from Harris, 1963; Koppitz, 1984). The training drawings are displayed in the top half of Fig. 1. For the first part of the experiment, the training 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 W&M, the training set consisted of two categories of six drawings, as shown in the top half of Fig. 1. At an abstract level, the first set (category A) were deemed to be relatively detailed and the second set (category B) were deemed to be relatively unusual (as confirmed by pilot testing reported by W&M); subjects learning that category A drawings were produced by ‘creative’ children would likely attend to the detail in the drawings. In addition to these abstract properties, drawings in the two categories could be distinguished on the basis of at least three simple rules based on perceptual features (W&M somewhat modified the drawings to attempt to conform to these rules). For category A, the drawings could be classified using simple rules like ‘curly hair and arms not at the sides’, or ‘light-colored shoes or smiling’, or ‘wearing a collar or tie’. For category B, the drawings could be classified using simple rules such as ‘straight hair or arms at the sides’, or ‘dark-colored shoes and not smiling’, or ‘ears and short sleeves’. Certainly, other simple rules based on other concrete perceptual features may also be available.

Sixteen transfer drawings, shown in the bottom half of Fig. 1, were selected and modified by W&M to attempt to combine abstract properties and concrete perceptual features of the training drawings in systematic ways. The first set of four drawings were not detailed and not unusual (T1–T4); the second set of four drawings were detailed and unusual (T5–T8); the third set of four drawings were detailed and not unusual (T9–T12); the fourth set of four drawings were not detailed and unusual (T13–T16). For each set of four drawings, two were best described by the concrete rules of category A and the other two were best described by the concrete rules of category B. For example, T1 and T2 had the concrete rules of category B, whereas T3 and T4 had the concrete rules of category A.

2.3. Procedure

Half of the subjects were instructed that they would be studying drawings by ‘creative’ or ‘non-creative’ children, while the other half 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 meaningful labels of ‘creative’ or ‘non-creative’, or meaningless labels of ‘group 1’ or ‘group 2’). Following W&M, as part of the study procedure, subjects were instructed to develop a written set of rules that could be used to divide up these drawings, as well as any possible new drawings, into the two categories. W&M systematically varied the assignment of category labels to sets of drawings. Because they reported that the specific assignment of labels to categories did not systematically influence the general conclusions of their study, we chose to have the ‘creative’ and ‘group 1’ labels apply to category A drawings and to have the ‘non-creative’ and ‘group 2’ labels apply to category B drawings.

In the next phase of the experiment, subjects were shown each of the training drawings one at a time on the computer, were asked to categorize them as ‘creative’

versus 'non-creative' or 'group 1' versus 'group 2' without any time pressure, were supplied corrective feedback, and were permitted to modify their written rules if necessary. The purpose of this phase of the experiment was to familiarize subjects with computer-generated versions of the drawings and to provide some practice with sequentially classifying the drawings into the two categories. After categorizing each of the drawings twice, their written rules were removed. The order of drawings was randomized for every subject.

In the next phase of the experiment, subjects were asked to categorize the training drawings and new transfer drawings. They were also instructed that they would receive no corrective feedback following their response. Half of the subjects were allowed to categorize the drawings without any time limit (unspeded condition), and the other half were instructed that they would be required to make their responses according to a response signal (speded condition).

In the speeded condition, a tone was presented 200 ms after the onset of the drawing. Subjects were required to make their categorization response within 300 ms 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 'non-creative' or 'group 1' or 'group 2') was displayed at the center of the screen. After 200 ms, the tone sounded and subjects were required to press the key associated with the displayed category label. If a subject responded before the tone sounded, he or she was informed to wait until the tone sounded; if a subject responded more than 300 ms after the tone sounded, he or she was informed to respond more quickly; if a subject made an incorrect response, he or she was informed not to make errors. The practice trials continued until the subject made 15 valid responses in a row. Once a subject had achieved this criterion, he or she was moved on to the test phase.

During the test phase, a crosshairs appeared at the center of the screen for 1 s. Then one of the 28 drawings (12 training and 16 transfer) was displayed and the subject was required to categorize it as 'creative' versus 'non-creative' or 'group 1' versus 'group 2'. In the unspeded condition, the subject could take as much time as needed. In the speeded condition, a tone sounded 200 ms after the drawing was displayed and subjects were required to make their response within 300 ms of the tone; subjects were informed if they had made a response prior to the tone or responded more than 300 ms after the tone. In both the speeded and unspeded conditions, no other corrective feedback was provided. After the response was made, the drawing was erased from the screen. After a 1 s 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 was randomly assigned to one of four groups: meaningful labels ('creative' versus 'non-creative') with unspeded responses, meaningful labels with speeded responses, meaningless labels ('group 1' versus 'group 2') with unspeded responses, and meaningless labels with speeded responses.




Fig. 2. Probability of classifying each of the 16 transfer drawings into category A, $P(A)$, as a function of label (meaningless labels in the left panel, meaningful labels in the right panel) and as a function of response deadline (hatched bars for unspeeeded, open bars for speeded). The column of italicized As and Bs on the right-hand side of each panel specifies how each drawing should be categorized if subjects appropriately used simple perceptual rules when provided meaningless labels (left panel) or if subjects appropriately used abstract properties when provided meaningful labels (right panel).

3. Results and discussion

The primary data of interest were the probabilities of classifying each stimulus presented during the testing phase as a member of category A, as a function of category label (meaningless versus meaningful) and as a function of response deadline (speeded versus unspeeeded). Fig. 2 displays the probability of classifying each of the new transfer drawings into category A, plotting the two category label conditions separately. In the speeded condition, categorization responses that were made prior to the response signal or more than 300 ms 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.

Recall that all subjects learned the same items in the same categories, and so both groups had equal opportunity to learn any abstract or concrete rules that existed. If the category labels had no influence on how subjects learned the categories, there should be no difference between the two groups provided the different labels. To the contrary, our findings in the unspeeeded condition replicate the important qualitative

results reported by W&M, 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 classified can best be explained by assuming that subjects provided meaningful labels tended to categorize on the basis of abstract properties, whereas subjects provided meaningless labels tended to categorize on the basis of concrete perceptual features. One way of quantifying this phenomenon is to define the correct category for each transfer drawing in terms of concrete perceptual features for subjects provided meaningless labels or in terms of abstract properties for subjects provided meaningful labels. These expected category responses are shown as columns of italicized As and Bs in both panels of Fig. 2. Focusing on the unspeeded condition, if we calculate accuracy in terms of these expected category responses, we find 70.8% performance by subjects provided meaningful labels and 61.6% performance by subjects provided meaningless labels (in the ensuing analysis we call this the appropriate mapping). As a contrast, we can instead calculate accuracy in terms of the expected category response for the *other* category label condition, calculating accuracy for subjects provided meaningless labels with respect to the expected category response for subjects provided meaningful labels and vice versa. When we do this, we find 49.1% performance for subjects provided meaningful labels and 34.8% performance for subjects provided meaningless labels (in the ensuing analysis we call this the inappropriate mapping). These performance differences were analyzed with a 2 (label) \times 2 (mapping) mixed analysis of variance with label (meaningless versus meaningful) as a between-subjects variable and mapping (appropriate versus inappropriate) as a within-subject variable. The importance of using the appropriate mapping to define accuracy was reflected by a significant main effects of mapping ($F(1, 28) = 23.48$), and an overall difference between meaningful and meaningless labels was reflected by a significant main effect of label ($F(1, 28) = 10.57$). No label \times mapping interaction was observed. An alpha level of 0.05 was established for all statistical tests.

The novel aspect of the present work is a finding that response deadlines had quite different influences on categorization decisions depending on whether meaningless or meaningful labels were provided. To see this, compare categorization probabilities under speeded and unspeeded conditions for subjects given meaningless labels, shown in the left panel of Fig. 2, with those for subjects given meaningful labels, shown in the right panel of Fig. 2. One way of summarizing these data is as follows. With meaningless labels, average 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. In other words, subjects given meaningless labels were much less likely to respond in a manner consistent with the perceptual features of the drawings under speeded conditions. By contrast, with meaningful labels, categorization probabilities actually tended to be somewhat further away from 50% in the speeded condition compared to the unspeeded condition. In other words, subjects given meaningful labels were somewhat more likely to respond in a manner consistent with the abstract properties of the drawings under speeded conditions. These results are further quantified in the following analysis.

A useful way of summarizing the effects of response deadlines requires slight

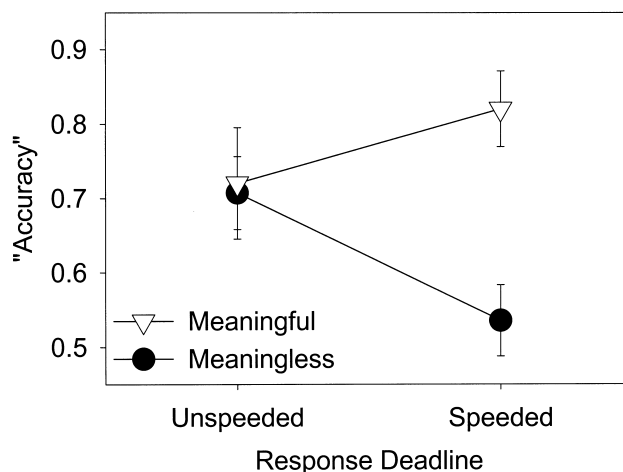


Fig. 3. Categorization 'accuracy' averaged across the transfer drawing as a function of response deadline (x-axis) and as a function of label (closed circles for meaningless labels, open triangles for meaningful labels). 'Accuracy' was defined with respect to the modal categorization response given to each stimulus in the unspeeded condition for each label type (see text).

recoding of the categorization responses as follows. As another measure of categorization 'accuracy' for new transfer drawings, for each stimulus, we simply defined the modal categorization response in the unspeeded condition as the 'correct' category response for that stimulus (e.g. for subjects given meaningless labels, the 'correct' response for T5 is category B, but for subjects given the meaningful labels, the 'correct' response for T5 is category A). For each stimulus, we then recoded each subject's category responses with respect to this nominally 'correct' response. Fig. 3 displays categorization 'accuracy' averaged across all the transfer stimuli. These accuracy data were analyzed with a 2 (label) \times 2 (deadline) between-subjects analysis of variance. The greater 'accuracy' for meaningful labels than meaningless labels was reflected by a significant main effect of label ($F(1, 56) = 7.33$). The differential influence of response deadline on the accuracy difference between the two types of label conditions was reflected by a significant label \times deadline interaction ($F(1, 56) = 6.12$). Planned comparisons revealed that with meaningless labels, subjects were significantly less 'accurate' in the speeded condition than the unspeeded condition (17.2% 'worse'); with meaningful labels, there was no significant difference between speeded and unspeeded conditions, although the trend was in the opposite direction (10.0% 'better').

A simple explanation for this difference in the effect of response deadline solely in terms of the amount of time needed to apply the rules in the two conditions seems unlikely. First, W&M systematically analyzed rules that subjects generated as a function of the kind of label provided and found that subjects given meaningful labels generated rules that were at least as complex as those generated by subjects given meaningless labels; a less systematic examination of the rules generated by our subjects provided similar results. Second, in the unspeeded conditions of the

present experiment, there was no significant difference in response times for subjects given meaningless (1420 ms) versus meaningful labels (1601 ms), suggesting that it does not necessarily take more time for subjects provided meaningless labels to apply the categorization rules they had generated.

That said, perhaps one reason why subjects provided meaningless labels were significantly disrupted by the speeded task is that their categorization rules require some kind of serial evaluation (e.g. checking the position of the arms, then checking the hair, and then checking the collar) that was disrupted by the requirement of making a rapid response (e.g. see Palmeri & Nosofsky, 1995; see also Smith & Sloman, 1994). By contrast, although rules generated by subjects provided meaningful labels were also quite 'complex' (see W&M), perhaps because these rules had a hierarchical or holistic nature to them (e.g. type of clothing or type of facial expression), they could be sufficiently evaluated within a limited amount of time. Another possibility, to be explored later, is that subjects provided meaningful labels may not need to make recourse to explicit rules at all, but could fall back on some kind of rapid non-analytic categorization process instead. Whatever the explanation, what is most striking is that the response deadline manipulation influenced categorization by subjects given meaningless labels but had less influence on categorization by subjects given meaningful labels, contrary to what some accounts of the effects of background knowledge may have predicted.

To summarize our results, we found that introducing a very rapid response deadline, causing subjects to make complex categorization responses in an average of just 370 ms after stimulus onset, had surprisingly little effect on performance when subjects were provided meaningful labels that tapped into a particular kind of background knowledge. It appears that whatever influence background knowledge has had on the category representations that are formed during learning or has had on the categorization strategies that are employed during classification of new drawings, that influence can take place at relatively early stages of a categorization decision (see also Lin & Murphy, 1997). The use of background knowledge need not be equated with a slow, optional, rule-based process that operates relatively late within the time-course of a categorization decision (e.g. Smith & Sloman, 1994). This seems consistent with the hypothesis that the joint influences of background knowledge and empirical evidence may sometimes be tightly coupled in the service of categorization (see Wisniewski, 1995). Important work remains to be done to more completely determine what specific influences background knowledge is having on the early stages of categorization.

Again, our results most clearly rule out the possibility that applying background knowledge is always a slow effortful process that comes online late within a categorization decision. Consistent with the kinds of rules subjects generated (see also W&M), perhaps the meaningful labels caused subjects to attend to global properties of the drawings (such as the amount of detail or the kinds of emotional expressions) that can be rapidly extracted when required by time limits (see Oliva & Schyns, 1997). Or, perhaps the background knowledge used by subjects in our experiment taps into previous instances of creative or uncreative drawings that subjects have experienced (e.g. Heit, 1994, 1998), and this instance information may be somehow

used for classifying drawings when time limits prohibit any use of the explicit rules subjects may have generated (see Logan, 1988; Palmeri, 1997; Palmeri & Nosofsky, 1995). We do not claim that using background knowledge will always lead to good performance in speeded conditions, but that at least in situations in which the form of background knowledge has a clear perceptual basis (creative versus non-creative drawings) it may indeed have a rapid influence. For other situations in which knowledge takes the form of complex causal networks (e.g. Ahn & Kim, in press) and in which stimuli are verbal feature descriptions (e.g. Smith & Sloman, 1994), it remains to be determined whether such background knowledge can have an influence during speeded conditions.

Acknowledgements

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 Grant BCS-9910756 from the National Science Foundation, Grant 1R01MH61370-01 from the National Institute of Mental Health, and a Grant from the Vanderbilt University Research Council awarded to the first author. We thank Woo-kyoung Ahn, Marci Flanery, Isabel Gauthier, Evan Heit, Douglas Medin, and two anonymous reviewers for their comments on earlier versions of this article.

References

- Ahn, W., & Kim, N. S. The causal status effect in categorization: an overview. In D. L. Medin (Ed.), *Psychology of learning and motivation*. San Diego, CA: Academic Press.
- 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.
- Hintzman, D. L. (1986). "Schema abstraction" in a multiple-trace memory model. *Psychological Review*, *93*, 411–428.
- Homa, D. (1984). On the nature of categories. In G. H. Bower (Ed.), *Psychology of learning and motivation* (Vol. 18, pp. 49–94). San Diego, CA: Academic Press.
- 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.
- Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review*, *95*, 492–527.
- Medin, D. L., & Schaffer, M. M. (1978). A context theory of classification learning. *Psychological Review*, *85*, 207–238.
- 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: 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. (1998). A rule-plus-exception model for classify objects in continuous-dimension spaces. *Psychonomic Bulletin & Review*, 5, 345–369.
- Nosofsky, R. M., Palmeri, T. J., & McKinley, S. C. (1994). Rule-plus-exception model of classification learning. *Psychological Review*, 101, 53–79.
- Oliva, A., & Schyns, P. G. (1997). Coarse blobs or fine edges? Evidence that information diagnosticity changes the perception of complex visual stimuli. *Cognitive Psychology*, 34, 72–107.
- 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., & 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.
- Reed, S. K. (1972). Pattern recognition and categorization. *Cognitive Psychology*, 3, 382–407.
- Smith, E. E., & Sloman, S. A. (1994). Similarity- versus rule-based categorization. *Memory & Cognition*, 22, 377–386.
- 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. (1994). On the interaction of theory and data in concept learning. *Cognitive Science*, 18, 221–281.