

Theory-based categorization under speeded conditions

CHRISTIAN C. LUHMANN

Vanderbilt University, Nashville, Tennessee

WOO-KYOUNG AHN

Yale University, New Haven, Connecticut

and

THOMAS J. PALMERI

Vanderbilt University, Nashville, Tennessee

It is widely accepted that similarity influences rapid categorization, whereas theories can influence only more leisurely category judgments. In contrast, we argue that it is not the type of knowledge used that determines categorization speed, but rather the complexity of the categorization processes. In two experiments, participants learned four categories of items, each consisting of three causally related features. Participants gave more weight to cause features than to effect features, even under speeded response conditions. Furthermore, the time required to make judgments was equivalent, regardless of whether participants were using causal knowledge or base-rate information. We argue that both causal knowledge and base-rate information, once precompiled during learning, can be used at roughly the same speeds during categorization, thus demonstrating an important parallel between these two types of knowledge.

Previous work on categorization has frequently focused on how people use similarity (i.e., featural overlap between instances and category representations; e.g., Nosofsky, 1986; Rosch & Mervis, 1975) or naive theories (i.e., general knowledge about how the world works; e.g., Murphy & Medin, 1985) to make category judgments. Typically, similarity use has been characterized as fast, automatic, and generally primary, whereas theory use has been characterized as slow, deliberate, and capable of overriding similarity only under particularly reflective situations (e.g., McRae, 2005; Sloman, 1996; see Kahneman, 2003, for a more general summary of this division). However, recent evidence has suggested that theory-based categorization can be fast (e.g., Lin & Murphy, 1997; Palmeri & Blalock, 2000).

Understanding when and why theory-based categorization can be fast allows us to place constraints on its underlying mechanisms. Our main theoretical assumption is that differences in the speed of categorization are due to differences in the complexity of the underlying *processes*, rather than differences in the *type* of information (i.e., similarity vs. theory). Within this framework, the pres-

ent study examines whether the effect of theories can be observed during rapid categorization. In this introduction, we first review previous work that examined fast and slow categorization using the traditional distinction of similarity versus theory. We then present two possible ways in which theory-based feature weighting can take place, one of which suggests simple, and therefore fast, processing.

Slow and Fast Use of Similarity and Theory Information

Evidence for similarity's automaticity comes from work showing that similarity is often used in situations in which it should not be. For example, Gentner and Toupin (1986) showed that participants would often use nonoptimal problem-solving strategies when problems shared superficial properties. Egeth (1966) showed that properties of pairs of stimuli were more rapidly identified as matching if other, irrelevant properties also matched. Similarly, Brooks, Norman, and Allen (1991) showed that participants were strongly influenced by feature overlap even when given an explicit classification rule. In contrast, theory use has been characterized as slow and deliberate. For example, Smith and Sloman (1994) failed to replicate a previously reported demonstration of theory use (Rips, 1989) simply by requiring participants to respond as quickly as possible. Only when participants were instructed to talk aloud while categorizing did the original, theory-based results reemerge. Similarly, Baraff and Coley (2003), described in Coley, Shafto, Stepanova, & Baraff, (2005) found that, with category-based induction,

This project was supported by National Institute of Mental Health Grant R01 MH57737 to W.-K.A., and by National Institute of Mental Health Grant R01 MH61370 and National Science Foundation Grant BCS-9910756 to T.J.P. We thank David Noelle and Art Markman for suggesting the precompiled knowledge explanation. Correspondence concerning this article should be sent to C. C. Luhmann, 2 Hillhouse Avenue, New Haven, CT 06511 (e-mail: christian.luhmann@yale.edu).

experts tended to rely more on specific domain theories than did novices. However, when placed under time pressure, even experts relied more on taxonomic similarity.

At a more theoretical level, Sloman (1996) construed theory-based categorization as a type of rule-based reasoning. For instance, dolphins are categorized as mammals on the basis of a rule such as: "If X gives live birth, produces milk, and has fur, then X is a mammal." Sloman (1996) claimed that "applying rules is relatively complex and slow" (p. 8). Similarly, McRae (2005) contrasted the use of statistical correlations with the use of causal relations in categorization. He argued that causal knowledge "influences performance on slower off-line tasks" (p. 57). However, it is not clear why applying rules or causal knowledge should necessarily be slower than utilizing similarity. Furthermore, several studies have provided demonstrations against these claims.

For instance, Goldstone and Medin (1994) showed that so-called local matches (e.g., matching a banana with a yellow bowl) dominated the early computation of similarity, whereas relational matches (e.g., matching lemons in a red basket with apples in a yellow bowl) only began to influence judgments when longer deliberation was allowed (e.g., over periods longer than 2,000 msec). Using composite stimuli (e.g., a row of three squares), Love, Rouder, and Wisniewski (1999) demonstrated that judging whether global forms (e.g., the fact that the row was made of three identical shapes) matched preceded matching the elements that comprised the global forms (e.g., the fact that the individual shapes were squares). Thus, these findings suggest that similarity does not lead to uniformly fast decision making.

Conversely, theories can sometimes exert influence during rapid categorization. For example, Lin and Murphy (1997) taught participants about novel objects and their intended functions (i.e., theories about form and function). They found that functionally central parts (e.g., "used to grab the neck of animals") influenced categorization judgments more than did functionally irrelevant parts (e.g., "used to hang the tool"), even when participants were instructed to categorize objects as fast as they could. However, this study lacked a comparable similarity-based condition that could be used to determine whether theory-based categorization was as fast as similarity-based categorization. Nonetheless, the mean categorization response time in this experiment was over 2,000 msec, which was in the range of what Goldstone and Medin (1994) considered to be moderately fast similarity-based judgments.

Palmeri and Blalock (2000) demonstrated much faster theory-based categorization using a response deadline procedure, in an extension of a study by Wisniewski and Medin (1994). They first established that meaningful category labels (drawings by "creative children" vs. "noncreative children") resulted in far richer categorization rules (e.g., "much more attention was given to the clothing" for drawings by creative children) than did neutral category labels (drawings by "group A" vs. "group B"), presumably because of participants' theories about creative and

noncreative children. Most critically, meaningful category labels significantly influenced how new drawings were categorized, even when the drawings were shown for only 200 msec and participants had only 300 msec to respond. Unfortunately, it is unclear what processes were actually responsible for the rapid theory-based categorization during the Palmeri and Blalock task, making it difficult to decide what critical differences between the paradigms produce fast (e.g., Palmeri & Blalock, 2000) versus slow (e.g., Smith & Sloman, 1994) theory use.

Slow Versus Fast Processing of Theories Illustrated by Feature Weighting

To understand these conflicting results, we propose that the speed at which different influences can be observed is determined by the complexity of the underlying processes, rather than the type of utilized information (i.e., similarity vs. theories). (See the General Discussion section for further discussion of how this theoretical framework can explain previous conflicting results.) To illustrate this notion, we describe a well-established effect of causal knowledge (i.e., theories) on feature weighting, which is the focus of the present experiments, and explain how this phenomenon could be slow or fast.

The theory-based feature weighting phenomenon we wish to investigate is the causal status effect (Ahn, Kim, Lassaline, & Dennis, 2000), in which the features of an object that act as causes in a domain theory are judged to be more important than the features that act as effects. For example, Ahn et al. presented participants with novel categories and a list of features (denoted abstractly here as X, Y, and Z). Participants were told that X causes Y and that Y causes Z. It was found that participants rated items that were missing feature Z to be the best category members and items that were missing feature X to be the worst category members, with items missing feature Y falling in the middle. In previous demonstrations of the causal status effect (e.g., Ahn et al., 2000), participants were allowed as much time as they needed to make their category judgments. Thus, it is unclear whether the processes leading to the causal status effect were relatively fast or relatively slow. Both of these possibilities appear plausible.

The causal status effect might be slow, because it results from online causal reasoning processes that operate while an exemplar is being categorized. For instance, Rehder (2003a, 2003b) suggested that upon encountering each exemplar, categorizers compute the likelihood that a given category's causal structure would produce the observed exemplar. For example, consider a case in which a reasoner knows that members of a category generally have two causally related features. Now imagine how a reasoner might categorize an exemplar possessing the cause feature but missing the effect feature. The reasoner might begin by judging the general likelihood of a cause feature's occurring, and then estimate the likelihood that the intervening causal relationship would fail. Then, by integrating the results of these inferences, a membership judgment could be computed. Thus, this "categorization

as causal reasoning” process (Rehder, 2003a, p. 739) is akin to an act of problem solving. If such problem-solving processes require conscious deliberation, categorization based on causal information would likely be slow.¹

Alternatively, the causal status effect might be fast, because it results from processing that took place when the category was first learned. That is, people may assign feature weights based on causal relations while they are initially learning the category (see Murphy & Medin, 1985, for a similar suggestion). People may then simply retrieve this precompiled information when asked to categorize an exemplar. Using these precompiled feature weights would require less time than an online computation using causal knowledge, and yet classifications would still embody the learned causal knowledge. According to this account, how feature weights are used at the time of categorization would not depend on whether the feature weights were computed from statistical information (e.g., base rates of features) or from domain theories (e.g., causal relations). In this sense, the act of categorization itself would not be qualitatively different from traditional similarity-based processes based on weighted feature matching (see Johansen & Palmeri, 2002, for a similar suggestion regarding rule use).

The present experiments attempted to evaluate whether the causal status effect can be demonstrated when rapid categorizations are required, or whether the causal status effect is limited to those situations that allow slow, deliberative categorizations. To investigate how the causal status effect is influenced by the speed of categorization, we employed a methodology similar to that of Ahn et al. (2000), but we manipulated the amount of time participants were allowed when making their judgments. In Experiment 1, we manipulated speed via an instructional manipulation, and in Experiment 2 we utilized a response deadline procedure. Experiment 2 also featured comparisons of situations in which feature weights were determined either by information about causal status or by information about feature frequency.

EXPERIMENT 1

Method

Participants. Twenty-nine Vanderbilt University undergraduates participated for course credit.

Materials. The stimuli consisted of four fictional animals, each having three features (denoted here as X, Y, and Z). The Appendix shows a full set of stimulus materials (under “Causal” in Table A1). The participants were told that X causes Y and Y causes Z. A summary diagram was provided that depicted causal relationships with arrows. The features were chosen to facilitate intuitive causal connections between adjacent features (e.g., a small heart causes a low body temperature, presumably because of weak circulation).

Procedure. The experiment was run on Apple iMacs using RSVP (Williams & Tarr, n.d.). During the learning phase, the participants were first given the list of features, the causal relationships between them, and the summary diagram for each category. To encourage the participants to think about the features causally rather than as an ordered list, they were instructed to “write about how you think each feature causes the next.”

After seeing descriptions for all four animals, the participants then received six blocks of trials. On each trial, the name of one of the animals was presented along with a table containing all 12 of the four animals’ features. The participants were asked to select (using a mouse click) the 3 features belonging to the presented animal. These features had to be selected in the appropriate order (see below) to be counted as correct. Four tables were used, each consisting of a unique arrangement of the 12 features. The table used on each trial was chosen randomly. This measure was taken to discourage the participants from responding simply on the basis of their memory for locations in the table. After each trial, the participants received feedback about their performance, as well as a summary of the presented animal’s features and the causal relationships. Selecting the correct features in the correct order for the entire set of animals two times in a row, with a maximum of one error, permitted a participant to move on to the next block.

In the first two blocks, responses were unsped. In the last four blocks, the participants were told to respond as quickly as they could, and trials that exceeded 15 sec were counted as incorrect. This speeded element was added in an attempt to automatize the use of the novel causal background knowledge, thereby approximating real-life lay theories. In addition, the participants were asked to respond with the animals’ features in the forward causal order (e.g., X, Y, Z) on three of the six blocks, and in the backward order (e.g., Z, Y, X) on the other three, to prevent a feature from being rated as central simply because it was always presented first in the list. The order manipulation alternated across blocks, always beginning with a forward block.

Upon completion of the learning phase, the participants proceeded to the transfer task. They received exemplars missing a single typical feature and were asked to rate the likelihood that the exemplar belonged to its target category, using an 8-point scale (with 1 labeled *definitely unlikely* and 8 labeled *definitely likely*).

Trials began with the name of one of the animals (the target category) appearing on the screen for 1 sec before being replaced by a triad of features. The three features belonged to the target category, with one feature negated by appending the phrase *does not* and displayed in red to facilitate reading of stimuli.² The features were presented in a triangular arrangement, with the location of each feature in the triangle randomly determined for each trial to discourage reliance on spatial information.

The participants received four blocks of these trials, each consisting of the 12 items (four categories with three features each) presented four times, for a total of 48 trials presented in random order. In two of the blocks (speeded condition), the participants were instructed to answer “as quickly as possible while still remaining accurate” and were given an example of how medical professionals often had to make decisions that were both accurate and rapid. In the other two blocks (unsped condition), the participants were told to “take as much time as needed.” The two conditions alternated across blocks and were counterbalanced across participants.

Results and Discussion

The RTs in the speeded blocks ($M = 1,560$ msec) were significantly faster than the RTs in the unsped blocks ($M = 3,202$ msec) [$t(28) = 10.97, p < .0001$]. This effect verifies the influence of our instructional speed manipulation.

In the unsped condition, items missing the terminal effect feature (Z) were rated as more likely category members than those missing a cause feature (X). More importantly, this causal status effect persisted in the speeded condition, as is shown in Figure 1. A 2 (speed condition: speeded vs. unsped) \times 3 (item type: missing X vs.

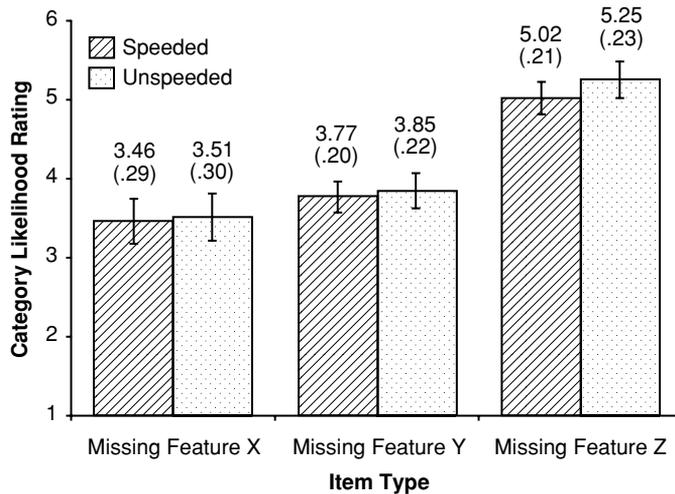


Figure 1. Results from Experiment 1. Error bars indicate \pm SEM. Means are presented with SEMs in parentheses.

missing Y vs. missing Z) repeated measures ANOVA revealed a significant main effect of item type [$F(2,56) = 22.69$, $MS_e = 44.39$, $p < .0001$], indicating that the three features did not influence judgments equally. Neither the main effect of speed [$F(2,56) = 2.89$, $p > .1$] nor the interaction between factors [$F(2,56) = 1.03$, $p > .3$] was significant, indicating that the causal status effect was not influenced by the speed manipulation. An analysis of the participants' first block of transfer trials demonstrated the same trend, ruling out practice effects as an explanation of speeded performance.

Planned comparisons were performed to further explore the effect of item type. Items missing feature Z were rated significantly higher than those missing the features X or Y in both the speeded [$t(28) = 5.00$, $p < .05$, and $t(28) = 5.68$, $p < .05$, respectively] and unspeeded [$t(28) = 5.07$, $p < .05$, and $t(28) = 5.25$, $p < .05$, respectively] conditions. Items missing feature Y were rated only marginally higher than those missing feature X in both the speeded [$t(28) = 1.47$, $p = .15$] and unspeeded [$t(28) = 1.47$, $p = .15$] conditions, possibly because the feature Y also served as a cause of another feature, making the difference between the X and Y features less pronounced. Overall, these results demonstrated that the causal status effect did not necessarily require prolonged deliberation (cf. Smith & Sloman, 1994). Indeed, when participants took only half the amount of time they normally would, the causal status effect was just as robust.

EXPERIMENTS 2A AND 2B

In Experiment 1, participants in the speeded condition were asked simply to respond as quickly as possible. Although this instructional manipulation allowed for naturalistic and spontaneous control of speed, the experiment did not allow us to assess the minimum time required to observe the causal status effect. In Experiments 2A and

2B, we imposed stricter control over participants' response times by enforcing specific response deadlines. Experiments 2A and 2B featured exactly the same method as Experiment 1; they differed only in the deadlines used. The deadlines used in Experiment 2A ranged from 5,000 msec (2,000 msec longer than the mean RT in Experiment 1's unspeeded condition) down to 750 msec. Experiment 2B used even shorter deadlines, ranging from 1,500 msec to 300 msec.

In addition, Experiment 2 allowed a comparison of the causal status effect with similarity-based feature weighting. Comparing the influence of causal and similarity information on categorization at each deadline would allow us to better understand the speed of the processes underlying the causal status effect. For example, the causal status effect may have been fast enough to withstand Experiment 1's instructional manipulation but still have been slower than the influence of similarity information. In addition, whereas Lin and Murphy (1997) and Palmeri and Blalock (2000) found fast theory-based categorization, both of those studies lacked a comparable similarity-based categorization condition. The results of Experiment 2 would resolve such ambiguities.

Earlier, we argued that similarity-based and theory-based categorizations could take place at roughly the same speed when the underlying processing complexity was equal. To make the processing complexity equal, we employed a comparable similarity-based condition, which utilized tasks as similar as possible to those used in our original theory-based categorization condition in Experiment 1, except that the feature weightings were determined through a similarity-based mechanism. We created a similarity-based determinant for feature weighting, by manipulating the frequency of each feature within a category (i.e., the percentage of category members that possessed a feature), because similarity is frequently calculated on the basis of how many attributes an item

has in common with other members of the category (see, e.g., Tversky, 1977). Indeed, base rates of features have been shown to correlate positively with typicality ratings (Gluck & Bower, 1988; Rosch & Mervis, 1975).

Experiments 2A and 2B utilized similarity conditions that provided frequency information in much the same way that causal information was provided in Experiment 1. As discussed earlier, we postulated that the causal status effect could occur when features were weighted by causal information *during learning*. If this should be the case, retrieving precompiled information about feature weighting would take the same amount of time *during categorization*, whether the weighting was determined by theories or similarity during learning.

Method

Participants. One hundred eight Vanderbilt University undergraduates were assigned to either the causal condition ($n = 25$ in Experiment 2A; $n = 30$ in Experiment 2B) or the base-rate condition ($n = 23$ in Experiment 2A; $n = 30$ in Experiment 2B).

Materials. The participants in the causal conditions were given the same stimuli and causal information that were used in Experiment 1. Those in the base-rate conditions were given the same stimuli, but all features were redistributed across the categories to discourage the spontaneous generation of causal relationships (see under Base-Rate in Table A1), and the causal information was replaced by the base rates of each feature. Thus, in the base-rate condition, each category was described as having three features (e.g., X, Y, and Z), with feature X appearing in 100% of category members, feature Y in 80%, and feature Z in 60%. If this experiment paralleled the effect of base rate on typicality judgments (Rosch & Mervis, 1975), then items in the base-rate condition that were missing feature Z (60%) would be rated as better category members than those missing feature X (100%), with feature Y (80%) falling somewhere in the middle.

Procedure. The learning phase for the causal conditions was identical to that of Experiment 1. That is, the participants received the descriptions of animals and causal relations, generated explanations for causal relations, and then completed a "selection task," in which they selected three features of a category in the requested causal order.

The learning phase for the base-rate conditions, designed to be similar to that for the causal conditions, also had three steps. The participants in the base-rate conditions first received the descriptions of animals one at a time, along with information that summarized the base rate for each feature. This step was similar to the first step in the causal conditions, in which those participants received the descriptions of animals and causal relations.

In the second step, the participants in the base-rate condition received a set of items constructed to mirror the stated base rates and were instructed to categorize each item into one of the stimulus categories. This procedure was adopted in place of writing explanations of the feature information, as the participants in the causal condition did, because such a task could conceivably have failed to reinforce the base-rate information while potentially reinforcing other, irrelevant (e.g., causal) information. Instead, the participants were given direct experience with the base rates on a trial-by-trial basis (as in, e.g., Koehler, 1996; Medin & Edelson, 1988).³ The participants received feedback after each trial, along with an additional summary of the animal's features and base rates. When a feature did not appear in an item, a feature from one of the other animals was substituted for the missing feature.⁴ Blocks of 25 trials were repeated until the participants achieved a categorization accuracy rate of 75% for any one block. Just as in the causal condition, this second step did not last long (e.g., 40% of the participants achieved it after one block; an additional 20%, after two blocks).

As the third and final step of the learning phase, the base-rate participants completed a "selection task" as in the causal condition, in which the participants matched features to categories in four blocks. The participants in the base-rate condition were instructed to select features in descending or ascending order dictated by their base rates (i.e., 100%–80%–60% or 60%–80%–100%). The feature matrix, exit conditions, speed manipulation, and order of blocks were all identical to those elements in the causal condition.

The transfer phase for both conditions was identical to the one used in Experiment 1, except for a modified speed manipulation. Instead of an instructional manipulation, Experiment 2 employed a signal-to-respond procedure (see, e.g., Lamberts, 1998). Thus, each feature triad was presented for a certain set amount of time ranging from 5,000 to 300 msec, after which the stimulus disappeared. The participants were instructed to respond to the item immediately after the item disappeared from the screen. Responses could not be made before the item disappeared or more than 300 msec after it had disappeared. If a response was not delivered within 300 msec of the stimulus's disappearance, the participants were told to respond more rapidly.

There were again four blocks of 48 trials during the transfer phase. Each block used one of four durations. In Experiment 2A, the durations were 5,000 msec (representing an "unsped" condition), 2,250 msec, 1,500 msec (the mean "sped" RT from Experiment 1), and 750 msec. In Experiment 2B, the durations used were 1,500, 750, 500, and 300 msec. Within each experiment, these blocks were ordered randomly for each participant.

Results and Discussion

The results from Experiment 2A are shown in Figure 2. The following analyses do not include 27.7% of the total responses, which were made either after the deadline or not at all. Note that the number of excluded responses did not vary as a function of either knowledge condition (26.0% in the base-rate condition vs. 29.3% in the causal condition) or speed condition (all F 's < 1).

To evaluate participants' responses in Experiment 2A, a 2 (knowledge condition: causal vs. base-rate) \times 4 (speed condition: 5,000 vs. 2,500 vs. 1,500 vs. 750 msec) \times 3 (item type: missing X vs. missing Y vs. missing Z) ANOVA was performed, with repeated measures on the latter two factors. The main effect of item type was significant [$F(2,90) = 64.77$, $MS_e = 205.27$, $p < .0001$], indicating that the three features did not influence judgments equally. This effect also significantly interacted with knowledge condition [$F(2,90) = 6.73$, $MS_e = 25.46$, $p < .005$], because the effect of item type was larger in the base-rate condition than in the causal condition. We also observed a significant main effect of speed [$F(3,135) = 3.11$, $MS_e = 2.07$, $p < .05$] and an interaction between item type and speed [$F(6,276) = 16.17$, $MS_e = 7.52$, $p < .0001$], because the effect of item type decreased as speed increased. The three-way interaction between item type, speed, and knowledge condition failed to reach significance (F 's < 1).

To explore the effect of item type, planned comparisons were conducted, broken down by each condition in Experiment 2A. The details of these comparisons, as well as descriptive statistics, can be found in Table A2. For both knowledge conditions, items missing feature Y always fell between the two other item types but were not consistently significant, probably as a result of lack of power. Examin-

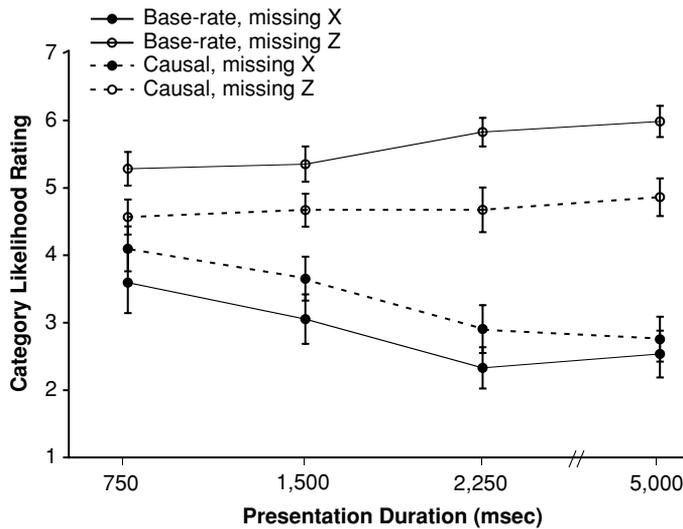


Figure 2. Results from Experiment 2A. Error bars indicate $\pm SEM$.

ing the difference predicted to be largest in both base-rate and causal conditions, the items missing feature Z were rated as significantly better category members than those items missing X at all deadlines in Experiment 2A (all $ps < .05$).

The results from Experiment 2B are shown in Figure 3. The following analyses do not include the 25.8% of responses that were not made by the deadline. The number of responses excluded did not vary as a function of knowledge condition (25.4% in the base-rate condition vs. 26.0% in the causal condition, $F < 1$), though it did vary with speed condition [$F(3,177) = 5.68$, $MS_e = 0.077$, $p < .005$]. This was due to the fact that fewer responses were excluded at the 300-msec deadline ($M = 20.4\%$ excluded) than at all other deadlines ($M = 27.5\%$ excluded, all $ts(59) > 2.79$, all $ps < .005$). Critically, this pattern held for both knowledge conditions, as demonstrated by a nonsignificant interaction between knowledge condition and speed condition [$F(3,174) < 1$].

To evaluate participants' responses in Experiment 2B, a 2 (knowledge condition: causal vs. base-rate) \times 4 (speed condition: 1,500 vs. 750 vs. 500 vs. 300 msec) \times 3 (item type: missing X vs. missing Y vs. missing Z) ANOVA was performed, with repeated measures on the latter two factors. We observed a significant main effect of item type [$F(2,114) = 9.46$, $MS_e = 8.60$, $p < .0005$] that did not interact with knowledge condition ($F < 1$), again indicating that the three features did not influence judgments equally. No significant main effect of speed was observed [$F(3,171) = 1.60$, $p = .19$], but this must be interpreted in light of a significant interaction between speed and item type [$F(6,342) = 4.04$, $MS_e = 1.64$, $p < .001$]. This interaction was again due to the fact that the effect of item type decreased as speed increased. This effect is further examined below. A significant main effect of background condition was observed [$F(1,57) = 5.22$, $MS_e = 67.42$,

$p < .05$], because judgments were slightly higher in the base-rate condition, but this factor did not interact with either of the other two factors. The three-way interaction between background condition, speed, and item type was also significant [$F(6,342) = 2.14$, $MS_e = 0.867$, $p < .05$], which is explained below. No other effects were significant (all $ps > .2$).

Planned comparisons were carried out for Experiment 2B to determine at which deadlines the background information had an effect on categorization. The details of these comparisons, as well as descriptive statistics, can be found in Table A2. For the base-rate condition, participants rated items missing feature X significantly lower than items missing feature Z in the 1,500- and 750-msec conditions, but not in the 500- and 300-msec conditions. In the causal condition, items missing feature X were rated significantly lower than those missing feature Z in the 1,500- and 750-msec conditions. This difference was marginally significant in the 500-msec condition ($p = .053$) and nonsignificant in the 300-msec condition. Ratings for items missing feature Y again fell between ratings for those missing feature X and those missing feature Z, but the differences were generally nonsignificant in both knowledge conditions (see Table A2).

In light of these further results, it is clear that the influence of causal knowledge can be observed under conditions significantly faster than those created in Experiment 1. Causal information continued to influence category judgments even when participants were allowed only 500 msec to view the exemplar and had to respond immediately afterward. Furthermore, although the influence of base-rate information on feature weighting was stronger than the effect of causal information at the slower speeds in Experiment 2A, we found no evidence that base-rate effects could be observed under faster conditions than could the effect of causal information.

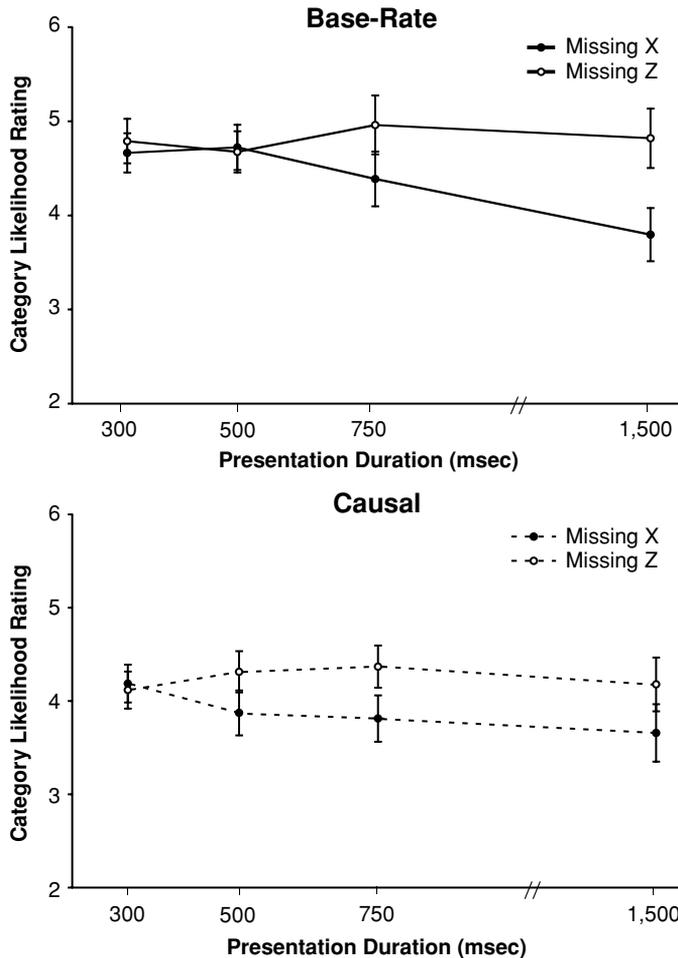


Figure 3. Results from Experiment 2B. The top graph contains the results from the base-rate condition. The bottom graph contains those from the causal condition. Error bars indicate $\pm SEM$.

GENERAL DISCUSSION

Summary

Previously, researchers have often argued that similarity-based effects should be observed when rapid judgments are required, whereas theory-based effects should only emerge once more leisurely judgments are allowed (McRae, 2005; Sloman, 1996). Our results suggest that the causal status effect, which describes how causal background knowledge influences feature weighting, is not limited to situations in which categorizers have ample opportunity to deliberate. Participants' category judgments exhibited the causal status effect even under speeded conditions. This was demonstrated with a naturalistic instructional manipulation, as well as with a more controlled response deadline manipulation. In fact, participants continued to exhibit the causal status effect when they were allowed only 500 msec to view the stimuli and only 300 msec to respond. In addition, causal and base-rate information elicited nearly identical patterns of behavior across all categorization speeds. These results

clearly contradict the prevailing notion that the effects of background knowledge can be found only in slow, deliberative settings.

Implications

These findings have several implications for models of categorization and category learning. First, the causal status effect was observed even during extremely rapid category judgments. This finding constrains the complexity of the processes operating at the point of categorization. Proposals that require significant cognitive processing to make category decisions may not provide general descriptions of how theory-based effects arise. In addition, proposals that posit the use of causal knowledge during the act of categorization itself (e.g., Rehder, 2003a, 2003b) must provide accounts of how such online computation can take place so rapidly.

Our results suggest that the speeded causal status effect is probably due to prestored feature weights established at the time of learning. If our participants were computing feature weights online and using causal knowledge

during each categorization, one would expect our causal condition to be noticeably slower than our base-rate condition, in which participants only retrieved stored base-rate information. Of course, it is possible that the deadlines used in the present study were too coarsely distributed to detect such a difference, but the marginal impact of causal knowledge at 500 msec, when similarity had no significant effect, suggests otherwise. What is more likely is that a simple feature weight computation was operating throughout the category learning process (see Lin & Murphy, 1997; Sloman, Love, & Ahn, 1998). As the category and its associated causal structure were learned, feature weights could be updated accordingly. Later, whenever a category decision was required, the feature weights would already be available to influence a relatively simple similarity computation.

Second, our findings suggest that the processes underlying the causal status effect (and perhaps other theory-based phenomena) may not always resemble the traditional conception of theory use. As previously mentioned, theory-based categorization has often been equated with deliberate rule use (Sloman, 1996). For example, bats and dolphins are mammals because they meet the criteria that define that class of animals. Smith and Sloman's (1994) demonstration of slow theory use likely involved classification rules that were both complex and novel (e.g., "If X is larger than 1 in. in diameter, then X is not a quarter"). In contrast to what the participants in our experiment had to do, Smith and Sloman's participants may have had to generate such rules (via background knowledge) before applying them to the observed exemplar. Thus, what Smith and Sloman interpreted as the influence of rule use may actually have been the influence of rule *generation* and/or *justification* in addition to the rule use itself. Thus, it seems reasonable that this set of processes is relatively slow.

Previous tests of speeded theory use are consistent with this interpretation. For example, Lin and Murphy's (1997) participants were told about an object's function and the role of each of the object's features in that function. The influence of theory use was demonstrated when participants weighted function-relevant features more heavily than function-irrelevant features (their actions appeared to conform to a rule such as: "If X has Y's function-relevant features, then X is a Y"). Such a rule has been posited to operate across the entire domain of artifacts (see, e.g., Bloom, 1996; Chaigneau, Barsalou, & Sloman, 2004) and thus could have influenced feature weighting relatively automatically. Similarly, Palmeri and Blalock (2000) asked participants to generate explicit classification rules before being tested with new drawings (in contrast to Smith & Sloman, 1994, in which participants presumably generated rules while attempting to perform the categorization). In the studies of Lin and Murphy as well as Palmeri and Blalock, the introduction of the categories, along with their associated theory-laden information (e.g., category labels), preceded the speeded categorization phase. This suggests that participants may have simply set the weights of the various features (e.g., function-relevant features)

ahead of time and then let the weights implement the influence of causal knowledge at the point of categorization (a possibility Lin & Murphy, 1997, acknowledge). If this is the case, there is no reason to think that the act of categorization should be slow.

Thus, Smith and Sloman (1994) may have mischaracterized their findings when they concluded that background knowledge can only influence deliberative acts of characterization. The critical factor might not have been the kind of information (i.e., world knowledge vs. similarity) that ultimately influences category judgments, but rather the complexity of the processes occurring at the point of categorization. Our findings also suggest that future models of categorization may need to reconsider how to integrate similarity information with background knowledge during learning and that the process of category learning itself may be able to transform one kind of knowledge into the other. The specific mechanisms responsible for transforming world knowledge into similarity information are currently unknown. Understanding these processes would greatly add to our understanding of conceptual representations and could act to unify seemingly disparate category information.

REFERENCES

- AHN, W.-K., KIM, N. S., LASSALINE, M. E., & DENNIS, M. J. (2000). Causal status as a determinant of feature centrality. *Cognitive Psychology*, **41**, 361-416.
- BARAFF, E., & COLEY, J. D. (2003, July). *Thinking about music: Novice and expert inductive reasoning*. Poster presented at the 25th Annual Conference of the Cognitive Science Society, Boston.
- BLOOM, P. (1996). Intention, history, and artifact concepts. *Cognition*, **60**, 1-29.
- BROOKS, L. R., NORMAN, G. R., & ALLEN, S. W. (1991). The role of specific similarity in a medical diagnostic task. *Journal of Experimental Psychology: General*, **120**, 278-287.
- CHAIGNEAU, S. E., BARSALOU, L. W., & SLOMAN, S. A. (2004). Assessing the causal structure of function. *Journal of Experimental Psychology: General*, **133**, 601-625.
- COLEY, J., SHAFTO, P., STEPANOVA, O., & BARAFF, E. (2005). Knowledge and category-based induction. In W.-K. Ahn, R. L. Goldstone, B. C. Love, A. B. Markman, & P. Wolff (Eds.), *Categorization inside and outside the laboratory: Essays in honor of Douglas L. Medin* (pp. 69-86). Washington, DC: American Psychological Association.
- EGETH, H. E. (1966). Parallel versus serial processes in multidimensional stimulus discrimination. *Perception & Psychophysics*, **1**, 245-252.
- GENTNER, D., & TOUPIN, C. (1986). Systematicity and surface similarity in the development of analogy. *Cognitive Science*, **10**, 277-300.
- GLUCK, M. A., & BOWER, G. H. (1988). From conditioning to category learning: An adaptive network model. *Journal of Experimental Psychology: General*, **117**, 227-247.
- GOLDSTONE, R. L., & MEDIN, D. L. (1994). Time course of comparison. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **20**, 29-50.
- JOHANSEN, M. K., & PALMERI, T. J. (2002). Are there representational shifts during category learning? *Cognitive Psychology*, **45**, 482-553.
- KAHNEMAN, D. (2003). A perspective on judgment and choice: Mapping bounded rationality. *American Psychologist*, **58**, 697-720.
- KOEHLER, J. J. (1996). The base-rate fallacy reconsidered: Descriptive, normative, and methodological challenges. *Behavioral & Brain Sciences*, **19**, 1-53.
- LAMBERTS, K. (1998). The time course of categorization. *Journal of Experimental Psychology: Learning, Memory, & 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 & Performance*, **23**, 1153-1169.

- LOVE, B. C., ROUDER, J. N., & WISNIEWSKI, E. J. (1999). A structural account of global and local processing. *Cognitive Psychology*, **38**, 291-316.
- McRAE, K. (2005). Semantic memory: Some insights from feature-based connectionist attractor networks. In B. H. Ross (Ed.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 45, pp. 41-86). San Diego: Academic Press.
- MEDIN, D. L., & EDELSON, S. M. (1988). Problem structure and the use of base-rate information from experience. *Journal of Experimental Psychology: General*, **117**, 68-85.
- MEDIN, D. L., & SHOBN, E. J. (1988). Context and structure in conceptual combination. *Cognitive Psychology*, **20**, 158-190.
- MURPHY, G. L., & MEDIN, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, **92**, 289-316.
- NOSOFSKY, R. M. (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **14**, 510-520.
- PALMERI, T. J., & BLALOCK, C. (2000). The role of background knowledge in speeded perceptual categorization. *Cognition*, **77**, B45-B57.
- REHDER, B. (2003a). Categorization as causal reasoning. *Cognitive Science*, **27**, 709-748.
- REHDER, B. (2003b). A causal-model theory of conceptual representation and categorization. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, **29**, 1141-1159.
- RIPS, L. J. (1989). Similarity, typicality, and categorization. In S. Vosniadou & A. Ortony (Eds.), *Similarity and analogical reasoning* (pp. 21-59). New York: Cambridge University Press.
- ROSCH, E., & MERVIS, C. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, **7**, 573-605.
- SLOMAN, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, **119**, 3-22.
- SLOMAN, S. A., LOVE, B. C., & AHN, W. (1998). Feature centrality and conceptual coherence. *Cognitive Science*, **22**, 189-228.
- SMITH, E. E., & SLOMAN, S. A. (1994). Similarity- versus rule-based categorization. *Memory & Cognition*, **22**, 377-386.
- TVERSKY, A. (1977). Features of similarity. *Psychological Review*, **84**, 327-352.
- WILLIAMS, P., & TARR, M. J. (n.d.). RSVP: Experimental control software for Mac OS [Computer software]. Retrieved July 10, 2002, from www.tarrlab.org/.
- WISNIEWSKI, E. J., & MEDIN, D. L. (1994). On the interaction of theory and data in concept learning. *Cognitive Science*, **18**, 221-281.

NOTES

1. Rehder's model does not explicitly make any predictions about the speed of categorization. It is described here as an example of a relatively sophisticated categorization process that could explain the influence of causal knowledge.

2. This should not have altered performance qualitatively, since the other two features were logically inferable from the negated feature, given that each item would contain the target category's features with exactly one feature negated. This task is equivalent to previously utilized feature weighting measures (e.g., Medin & Shoben, 1988; Sloman, Love, & Ahn, 1998) in which participants were asked whether an object, which resembled a target category (e.g., a dog) in all ways except that it did not have a target feature (e.g., four legs), was a member of a target category.

3. The additional training with specific exemplars differed from the causal condition. Nonetheless, the increased exposure to the base rates provided a stronger test of the fast theory-based categorization hypothesis.

4. When a feature was missing, it was replaced and not negated. This was done because items in the transfer phase used negation, and we wanted participants to have neither practice nor feedback on the transfer items. Arguably, this might have offered a more difficult learning situation for the base-rate condition because of increased similarity between categories, but because the exit condition for the learning phase was equated between the base-rate and causal conditions, the base-rate condition would have offered more training and thus provided a stronger test against our hypothesis.

APPENDIX

Table A1 below shows stimuli used in Experiments 1 and 2. In the causal condition, the three features of each category formed a causal chain in the order presented (e.g., the kehoes' small hearts cause them to have a low body temperature, which in turn causes them to hibernate during winter). In the base-rate condition, the features within a category were presented in descending order of base rates (e.g., 100% of kehoes have small hearts, 80% of kehoes have heavy fur, and 60% of kehoes eat aquatic plants).

To ensure equal salience of the three features in each category in the absence of causal information, a separate group of 29 participants received the list features without causal or feature frequency information and rated the likelihood of category membership for items missing a single feature. A 2 (knowledge condition: causal vs. base-rate) \times 3 (item type: missing X vs. missing Y vs. missing Z) ANOVA with repeated measures on both factors showed no significant effect of knowledge condition [$F(1,28) = 1.51, p > .3$] or item type ($F < 1$) and no significant interaction between these factors ($F < 1$).

Although the pretest results reported in Experiment 1 demonstrated no a priori difference in feature salience without causal information, an extra measure was taken in developing the base-rate condition stimulus materials in order to equate the two conditions. Thus, for instance, features that served as initial causes in the causal condition were always the 100% features in the base-rate condition, and so on.

Table A1
Conditions and Stimuli in Experiments 1 and 2

Causal	Base Rate
Kehoes	Kehoes
have small hearts	have small hearts
have a low body temp	have heavy fur
hibernate during winter	eat aquatic plants
Javans	Javans
have small stomachs	have small stomachs
eat often	have a low body temp
sleep for short periods	groom their peers
Ludaks	Ludaks
have large lungs	have large lungs
hold breath well	eat often
eat aquatic plants	hibernate during winter
Andals	Andals
live in cold climates	live in cold climates
have heavy fur	hold breath well
groom their peers	sleep for short periods

Table A2
Statistics

Experiment	Condition	RT (msec)	Descriptive Statistics						<i>t</i> Tests					
			Missing X		Missing Y		Missing Z		Missing X vs. Missing Y		Missing Y vs. Missing Z		Missing X vs. Missing Z	
			<i>M</i>	<i>SE</i>	<i>M</i>	<i>SE</i>	<i>M</i>	<i>SE</i>	<i>t</i>	<i>p</i>	<i>t</i>	<i>p</i>	<i>t</i>	<i>p</i>
2A	Causal	5,000	2.89	0.33	4.43	0.31	4.92	0.28	4.37	**	2.60	*	5.84	**
		2,250	3.03	0.34	4.48	0.30	4.70	0.32	3.90	**	1.21	n.s.	4.02	**
		1,500	3.74	0.31	4.43	0.26	4.59	0.24	2.73	*	1.21	n.s.	3.19	**
		750	4.08	0.31	4.51	0.29	4.57	0.27	1.92	n.s.	0.66	n.s.	2.18	*
	Base rate	5,000	2.57	0.34	5.12	0.23	5.97	0.23	7.17	**	4.47	**	7.68	**
		2,250	2.49	0.33	4.98	0.24	5.82	0.21	6.78	**	3.74	**	7.24	**
		1,500	3.19	0.36	5.10	0.24	5.33	0.25	5.41	**	1.33	n.s.	5.19	**
		750	3.59	0.43	5.11	0.28	5.30	0.25	4.38	**	1.04	n.s.	3.93	**
2B	Causal	1,500	3.66	0.29	3.90	0.27	4.18	0.31	1.55	n.s.	2.04	n.s.	2.22	*
		750	3.81	0.23	4.01	0.25	4.37	0.25	1.59	n.s.	2.15	*	2.76	**
		500	3.87	0.22	4.12	0.20	4.31	0.24	1.54	n.s.	1.18	n.s.	2.01	n.s.
		300	4.18	0.20	4.23	0.19	4.12	0.20	0.47	n.s.	1.33	n.s.	0.75	n.s.
	Base rate	1,500	3.80	0.32	4.69	0.30	4.82	0.28	3.15	*	0.64	n.s.	3.16	*
		750	4.39	0.31	4.82	0.28	4.96	0.29	2.00	n.s.	0.94	n.s.	2.05	*
		500	4.72	0.22	4.68	0.24	4.68	0.24	0.51	n.s.	0.92	n.s.	0.86	n.s.
		300	4.66	0.24	4.89	0.23	4.79	0.21	1.85	n.s.	0.84	n.s.	0.92	n.s.

Note—Means and standard errors are provided for each condition. For Experiment 2A, $df = 22$ in the t test for the base-rate condition, and 24 for the causal condition. For Experiment 2B, $df = 29$ for both conditions. n.s.: $p > .05$. * $p < .05$. ** $p < .01$.