Knowledge Effect and Selective Attention in Category Learning: An Eyetracking Study

ShinWoo Kim (shinwoo.kim@nyu.edu)
Bob Rehder (bob.rehder@nyu.edu)
Department of Psychology, New York University
6 Washington Place, New York, NY 10003 USA

Abstract

Two experiments tested the effect of prior knowledge on attention allocation in category learning. Using eyetracking, we found that (a) knowledge affects dimensional attention allocation, with knowledge-relevant features being fixated more often than irrelevant ones, (b) this effect was not due to initial attention bias to the relevant dimensions but rather gradually emerged in response to observing category members, and (c) the effect grew even after the last error trial, that is, in the absence of error. These results pose challenge to current models of knowledge-based category learning.

Because of the importance of categories for human cognition, the manner in which people learn categories has received intensive study. Among many procedures, supervised classification learning has been popular, and a number of basic facts have been established.

One of these concerns the role of selective attention in category learning. Selective attention has played a prominent role in theories of category learning since the finding that learning difficulty correlated with the number of diagnostic dimensions needed for classification (Shepard, Hovland, & Jenkins, 1961). In both exemplar models and prototype models, selective attention is explicitly formalized in terms of the weights of dimensions on classification (e.g., Kruschke, 1993; Nosofsky, 1992). Rule-based models also assume selective attention to the dimensions referred to by the current hypothesis (Smith, Patalano, & Jonides, 1998).

A second finding concerns the role of the prior knowledge. Although early research focused on the effect of empirical information, subsequent research has shown that theoretical knowledge influences every type of category-based judgment that has been tested, including learning, induction, and categorization (Murphy, 2002). For example, subjects usually learn categories far faster when a category's diagnostic features can be related to a common theme, and also show better learning of those related features versus unrelated ones (Murphy & Allopenna, 1994).

Although attention and knowledge each has been shown to be important, little is understood on how knowledge affects attention in category learning. This question is important because any theory that accounts for knowledge-based category learning (e.g., Heit & Bott, 2000; Rehder & Murphy, 2003) is incomplete in the absence of attention mechanism. However, this omission is understandable because virtually nothing is known about how knowledge affects attention: modeling attentional effects is impossible if there is no data to model.

Knowledge Effect on Selective Attention

Using eyetracking, we addressed three open questions regarding how knowledge affects dimensional attention during category learning. The first concerns whether in fact prior knowledge affects dimensional attention. If a subset of category features is knowledge-related, it is likely that such knowledge directs attention to the related features and away from unrelated ones (Kaplan & Murphy, 2000). However, knowledge might instead change how features are processed and encoded without changing what is attended. For example, knowledge might allow knowledge-related features to be associated more strongly with its underlying representation. Thus, our first goal is to determine whether knowledge indeed induces any change to what is attended.

Assuming that it does, the second question concerns the time course of that effect. Some theorists have suggested that one role of knowledge is to *preselect* dimensions for further testing. For example, Pazzani's (1991) rule-based PostHoc model selectively attends to goal-relevant features and thus predicts preselection of those dimensions. Kruschke (1993) suggested that his associative ALCOVE model can account for knowledge by setting *initial* attention weights on the related dimensions higher than on others.

However, knowledge effect on attention might emerge gradually in response to observing category members. Because prior knowledge consists of representations in the long-term memory, multiple exemplars may need to be observed for it to become sufficiently active. Or, a learner may only begin to make use of knowledge when a simpler strategy yields error signals. Thus, our second goal concerns whether preference for attending to related dimensions emerges gradually with experience of category members.

Assuming that it does, the third question concerns whether error feedback is required to mediate change in attention. One possibility is that (even) greater attention to related features occurs in the absence of error, because merely observing related features might be sufficient to further activate prior knowledge. However, supervised learning is often characterized as learners' adapting their responses to error feedback so as to minimize error (Kruschke, 2001). It is assumed that doing so requires adjusting attention to more diagnostic dimensions. This error-driven account might be extended for knowledge-based category learning, for example, error feedback might serve as cue indicating to the learner to use prior knowledge (and attention might shift as a result). Thus, our third question concerns whether such shifts occur only while learners are committing errors or even in the absence of error, that is, after solving the classification problem.

To address these questions, we conducted an eyetracking study of thematic category learning (e.g., Murphy & Allopenna, 1994). In recent years, eyetracking has been successfully applied to studying dimensional attention in various category learning tasks (Rehder & Hoffman, 2005ab; Rehder, Colner, & Hoffman, 2009). We now use eyetracking to study knowledge effect on attention.

Overview of the Experiments

Two categories of ants (i.e., Dax and Kez) were constructed from six binary dimensions using one-away structure (Table 1). Figure 1 illustrates two example prototypes. Four dimensions were described as useful in either a cold, tundra-like environment or a hot, desert-like environment. The other two "neutral" dimensions were unrelated to these themes. Table 2 presents example feature descriptions for the prototypes in Figure 1. To prevent the themes from being blatantly obvious, the "tundra" and "desert" themes were not mentioned, but only indirectly suggested by the feature descriptions.

In Experiment 1, we conducted a non-eyetracking study to establish whether our novel materials would induce standard behavioral effects of prior knowledge on category learning. In Experiment 2, we conducted an eyetracking study to address the three questions we raised above.

Experiment 1

Materials. Dax and Kez categories were constructed from six binary dimensions. Table 1 presents category structure in which the prototypes are 111111 (for Dax) and 000000 (for Kez). Four feature assignments to Dax/Kez prototypes were used to balance features to categories: 111111/000000, 101010/010101, 010101/101010, and 000000/111111.

In the related condition, Daxes were related to the tundra theme and Kezes to the desert. Of the six dimensions, four were related to the themes and the other two were neutral. The neutral dimensions were either tail/foot, wing/mouth, or forearm/antenna with the remainder being theme-related. In the unrelated condition, all features were neutral.

The two experimental conditions (related vs. unrelated), the four feature assignments to categories, and the three related/neutral dimension assignments resulted in 24 cells.

Table 1: Structures of Dax and Kez categories.

	Dimensions					
Exemplars	Tail	Foot	Wing	Mouth	Forearm	Antenna
Dax						
D0	1	1	1	1	1	1
D1	1	1	1	1	1	0
D2	1	1	1	1	0	1
D3	1	1	1	0	1	1
D4	1	1	0	1	1	1
D5	1	0	1	1	1	1
D6	0	1	1	1	1	1
Kez						
K0	0	0	0	0	0	0
K1	0	0	0	0	0	1
K2	0	0	0	0	1	0
K3	0	0	0	1	0	0
K4	0	0	1	0	0	0
K5	0	1	0	0	0	0
K6	1	0	0	0	0	0

Participants. Thirty NYU students were randomly assigned to the 24 cells with a constraint of at least one person in one cell (related, n = 14; unrelated, n = 16).

Procedure. The experiment consisted of three phases: knowledge acquisition, category learning, and a single-feature test. In knowledge acquisition, participants studied a total of 12 features, six from each category. Each screen displayed an ant with one visible feature and the other five features hidden behind gray rectangles. Below the ant were descriptions of the visible feature. Subjects studied the 12 features on their own pace by navigating 12 screens with left/right arrows keys.

To ensure learning, participants were required to take a multiple-choice test followed by a recall test. Both tests consisted of 12 questions, one for each feature. In the multiple-choice test, a question presented an ant with one visible feature (as in the learning screens), and participants chose one of the four alternatives below the ant. Immediate feedback was provided for each question, and after the test, total number of errors was also provided. When any error occurred, they were returned to the initial learning screens for additional study, and then retook the test that presented only the questions they missed. This process repeated until all questions were correctly answered. In the recall test, subjects verbally described the visible feature instead of making a choice. Immediate feedback was provided for each question, and after the test, total number of errors was provided. Any error during the recall test obligated the subject to restart the knowledge acquisition all over again.

The category learning phase began with two practice trials followed by training blocks that randomly presented 14 exemplars, seven from each category (Table 1). subjects classified each as either a Dax or Kez by pressing "z" or "/". Immediate feedback was provided below the exemplar ("Correct" or "Wrong") and the exemplar remained visible for 3.8 s after the response. The training ended after two consecutive errorless blocks or after the 15th block.

Finally, a single-feature test followed category learning. Each trial presented an ant displaying one visible feature (as in the learning screens), and subjects classified the visible feature (as in training). No feedback was provided. After each choice, subjects rated confidence by positioning a slider on a scale, which was then converted to [0–100] range.

Results

Subjects were very accurate in the test during knowledge acquisition. 22 subjects committed zero errors. Related (.97) and unrelated (.98) subjects were equally accurate, $t < 1$.

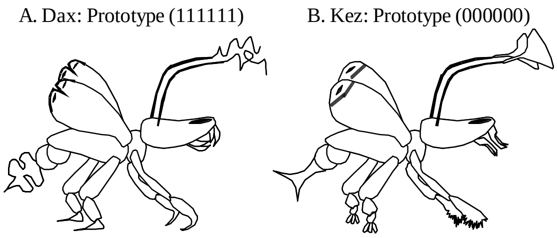

Figure 1: Two categories' example prototypes.

Table 2: Example feature descriptions.

Dimension	Dax [Tundra Theme]	Kez [Desert Theme]
Related		
Antenna	Because the temperature is very low, parts of ants' eyes (e.g., cornea, iris, pupil) often freeze and the ants become blind. When that happens, this thread type of flexible antennae is used to detect close objects.	Because the air is hot and dry, the ants are vulnerable to dehydration. To maintain hydration, the ants use this fan type of antennae to absorb water vapor from the air.
Mouth	Because the ground is frozen, the ants need to cut and break tough soil in search of their food. This type of mouth with sharp incisors serves this function.	Because sources of food are covered with sand, they need to be cleared before swallowing. The inner surface of the ants' mouth has short but stiff hairs that filter out these impurities.
Forearm	Because of frequent blizzards, the ants need to anchor themselves during high winds. This type of forearm allows the ant to hold its position.	Because the ants' prey (e.g., fleas) hide in sand, the ants use this type of forearm to sweep the sand and detect the prey.
Foot	Because the ground surface is slippery, the ants need to have wide feet to maintain their footing.	Because the ground surface is extremely hot, the ants switch the toe that comes into contact with the ground in each step to avoid burning.
Neutral		
Tail	The ants feed proteins stored in the humps to their larvae using the sharp nozzle in the end of tail.	The ants lay a large number of eggs at a time. This trumpet-shaped tail allows the ants to deliver a large number of eggs.
Wings	While flying, the ants control their rapid changes in direction by adjusting the fore- and rear-flaps in each wing.	The ants have red spots in the wing ends. The color becomes brighter in the mating season by the hormones produced in the gray area.

During training, 12 related and 14 unrelated subjects reached the learning criterion of two errorless blocks. The related learners did so in fewer blocks (5.50) than the unrelated learners (8.57), $t(24) = 2.85$, $p < .01$, while committing fewer total errors (8.67 vs. 19.07), $t(24) = 3.29$, $p < .01$. These results replicate standard behavioral effect of faster learning in thematic categories.

Single-feature test (Table 3) also replicated greater accuracy on related dimensions (.89) than on the neutral ones (.71), $t(11) = 1.79$, $p = .05$ (one-tailed) and than on the neutral dimensions in the unrelated condition (.76), $t(24) = 1.85$, $p = .07$. The more sensitive signed confidence ratings were computed by negating the ratings in the incorrect trials (Table 3). Related learners' signed ratings were greater than 0 for both dimension types, p's < .01. More importantly and consistent with the accuracy, the ratings were greater for the related dimensions (67.1) than for the neutral ones (37.0), $t(11) = 1.82$, $p < .05$, and for the neutral dimensions in the unrelated condition (43.8), $t(24) = 2.32$, $p < .05$.

Table 3:Single -feature test results .

	Related		Unrelated
	Related Dimensions	Neutral Dimensions	Neutral Dimensions
Experiment 1			
Accuracy	.89	.71	.76
Signed rating	67.1	37.0	43.8
Experiment 2			
Accuracy	.91	.70	
Signed rating	73.6	29.1	

Discussion

Experiment 1 confirmed the standard learning advantage of thematic categories with our novel materials. Related learners learned to distinguish the categories faster with fewer number of errors, and showed better single-feature test performance on the related features than on the neutral ones. The 5.5 blocks of learning speed suggests that the categories were not pre-learned in the knowledge acquisition phase.

Experiment 2

To address the main questions, we set out to measure eye movements during training. Because unrelated condition was included as a control for learning performance, we only tested the related condition in Experiment 2.

Materials. The materials were the same as in Experiment 1.

Participants. Twenty-four NYU students were randomly assigned in equal numbers to one of the four assignments of features to categories and to one of the three assignments of related/neutral dimensions.

Procedure. The procedure was the same as in Experiment 1, with a few additional steps for eyetracking during training. Participants were first fitted and calibrated to the eyetracker. Each trial began with a drift correction that compensated small movements of the eyetracker on the head. We used a gaze-contingent display such that a feature was fully visible when it was fixated but blurred when it was not, to minimize use of peripheral vision. After each classification, auditory feedback indicated whether the response was correct (chime) or incorrect (ding).

Eyetracking measures. The eyetracker yields a stream of fixations and their corresponding x-y screen locations and durations. We defined six circular areas of interest (AOIs) that encompass the features on the monitor. All fixations outside of the AOIs were discarded, as were those that occurred after classification response. Using the remaining fixations, we computed four measures on each trial.

The first is the *number of dimensions observed* in each trial. We counted a dimension "observed" if that dimension is fixated at least once, and thus, it ranges [0–4] for the related dimensions and [0–2] for the neutral ones. The second, *fixation probability*, is obtained by dividing the number of dimensions observed by 4 and 2 for the related and neutral dimensions, respectively. The third, *proportion fixation number*, is computed by taking the number fixations to the related dimensions and dividing it by the total number of fixations. The fourth, *proportion fixation time*, is the result of taking the time fixating the related dimensions and dividing it by the total fixation time in each trial.

Results

Basic learning results. Once again, participants were very accurate in the multiple-choice and recall tests in knowledge acquisition (avg. accuracy = .97). During training, 20 of 24 participants reached the learning criterion of two errorless blocks (M = 6.5 blocks; cf., 5.5 in Expt. 1) while committing 10.60 total errors (cf., 8.67 in Expt. 1).

Single-feature test. As in Experiment 1 (Table 3), learners exhibited greater accuracy on the related dimensions (.91) than on the neutral ones (.70), $t(19) = 3.31$, $p < .01$. Once again, the signed confidence ratings were greater than 0 for both dimension types, p's < .01, and greater for the related dimensions (73.6) than for the neutral ones (29.1), $t(19) = 5.75$, $p < .001$.

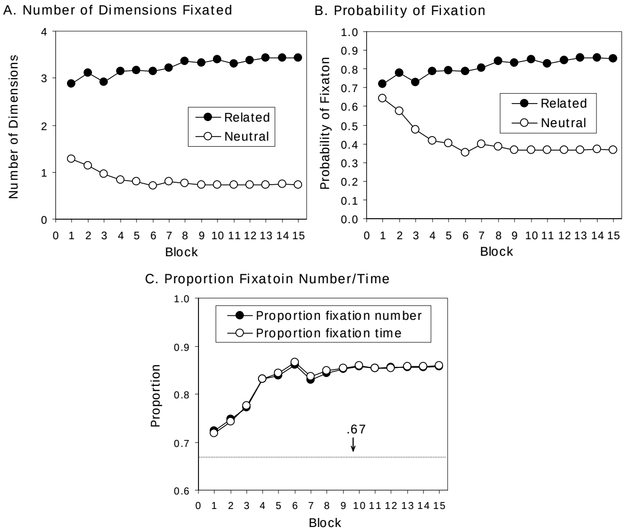

Figure 2: Eye fixation results from Experiment 2.

Eye fixations. Figure 2 reports learners' eye fixations averaged in each block. In construction of this figure, we assumed that the learners' eye movements in the remaining blocks after reaching criterion would have been identical to those in their last block.

Figure 2A shows that learners initially observed 3 of the 4 related dimensions and 1.25 of the 2 neutral dimensions. Figure 2B equates different number of dimension types. The figure indicates that learners fixated the two types of dimensions with about equal probability but become more (less) likely to fixate related (neutral) dimensions.

A 2 x 2 within-subjects ANOVA was conducted on fixation probabilities in Figure 2B with dimension type (related vs. neutral) and block (first vs. last) as factors. There was a main effect of dimension, $F(1, 19) = 20.294$, $p < .001$, indicating overall greater chance of fixating related dimensions. There was no main effect of block ($p > .10$), but a significant interaction between dimension type and block, $F(1, 19) = 25.904$, $p < .001$, suggested the increase (decrease) in fixating the related (neutral) dimensions. T-tests revealed that learners were more likely to fixate the related dimensions than the neutral ones in all blocks, p's < .03, except block 1, $p > .09$. The small (nonsignificant) difference in block 1 might have resulted from fixations in the later trials of the block. The nearly identical probabilities for the related (.73) and neutral (.75) dimensions on the first trial suggest absence of strong initial preference.

These results are further supported by more sensitive proportion measures in Figures 2C. Because there were four related and two neutral dimensions, a value of .67 (= 4/6) reflects bias toward neither dimensions. Both proportions start off slightly greater than .67 and then shift in favor of the related dimensions. T-tests comparing the first and last blocks confirmed increase in both proportions, p's < .001. In addition, both proportions were greater than .67 in all blocks except blocks 1 and 2, p's < .02. This result is consistent with the fixation probabilities in Figure 2B indicating no attentional *preselection* of the related dimensions (although very weak but nonsignificant initial preference is obtained).

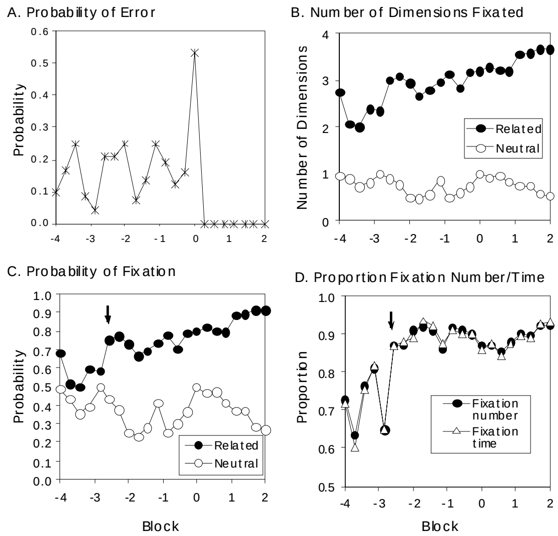

Figure 3: Backward learning curves from Experiment 2.

Backward learning curves. Figure 2 indicated learners' gradual shift in attention during training. We also asked how that shift relates to reduction in error. To address this question, we created backward learning curves (Figure 3) by translating each subjects' trial numbers so that their last error occurred on trial 0. Because the primary interest was knowledge use and its relation to error reduction, we included only 14 of the 20 learners (i.e., knowledge users) whose eye movements and single-feature test results both showed obvious bias for the related dimensions. In Figure 3, we included 60 trials (about 4 blocks) before the last error and 28 trials (the 2 blocks of learning criterion) after the last error. We averaged every 4 trials to obtain one data point.

Figure 3A presents the probability of error that indicates nonzero but small number of errors until the last error. Of greater interest are Figures 3B (number of dimensions fixated) and 3C (fixation probabilities). First, consider the eye fixations before the last error. Both figures indicate that shift in attention begins from about three blocks before the last error. T-tests confirmed that the related dimension were fixated with greater probability than the neutral ones starting from the data point indicated by an arrow (Figure 3C), p's < .05. The proportion measures (Figure 3D) were also reliably greater than .67 from the same data point, p's < .01. These results establish that the knowledge users began shift in attention well before solving the classification problem.

Next consider the eye fixations after the last error. Figures 3B, 3C, and 3D all indicate continued shift in attention after the last error, that is, despite in the absence of error feedback. A 2 x 7 within-subjects ANOVA was conducted on the fixation probabilities (Figure 3C) with dimension type (related vs. neutral) and data point (1 to 7) as factors. There was a main effect of dimension type, $F(1, 13) = 40.617$, $p < .001$, confirming grater chance of fixating related dimensions. There was no main effect of data point ($F < 1$), but significant interaction, $F(6, 78) = 5.796$, $p < .001$, confirmed the increase (decrease) in fixating the related (neutral) dimensions. Considering the two dimension types separately, fixation probabilities (Figure 3C) increased from the first two data points to the last two ($p = .10$) for the related dimensions and decreased for the neutral dimensions ($p = .09$). Finally, the two proportions (Figure 3D) also showed reliable increase for the same data points, p's < .05.

Discussion

Experiment 2 addressed our three main questions. First, eye fixations showed that prior knowledge indeed affects dimensional attention allocation, as learners devoted more attention to related dimensions. Second, learners showed only a weak or no initial attentional preference but then gradually shifted attention to related dimensions. Third, this shift in attention continued after the classification problem was solved, that is, in the absence of error feedback.

General Discussion

Previous research has documented large knowledge effects on category learning, and many investigators considered the possibility that these effects are mediated by attention. We discuss our findings regarding the questions we raised and the implications they have on models of knowledge-based category learning.

An effect of prior knowledge on attention. Using eyetracking, we answered the most basic question that knowledge indeed affects dimensional attention. Contra the encoding account in the basic memory literature, we found a robust effect of knowledge on dimensional attention. To our knowledge, this is the first direct confirmation of the frequent proposals that knowledge directs attention to knowledge-relevant information (e.g., Heit & Bott, 2000; Kruschke, 1993; Murphy & Medin, 1985; Pazzani, 1991).

Regarding feature learning, both Baywatch (Heit & Bott, 2000) and KRES (Rehder & Murphy, 2003) correctly predicted better learning of related features (Experiments 1 & 2). However, in these models, this result has been accounted for not by attention allocation but solely by various forms of cue competition that arise from error-driven learning. For example, in Baywatch, related features are learned faster because they are additionally connected to the category label via common prior concept units that accelerate their learning at the cost of the neutral features.

However, it is well-known that many standard effect of cue competition can arise not only from dynamics of error-driven processes but also by attentional mechanisms (e.g., Kruschke, 2001). The present results suggest that such attentional effect on cue competition on feature learning also hold in knowledge-based category learning: Because learners attend the neutral features less often, they will be learned less well. Thus, neutral features were at a double disadvantage in *learning*, suffering from not only the effects of cue competition but also fewer attentional resources.

Regarding features' contribution to final classification decisions, the neutral features were learned less well means that they are contributing less to final classification decisions. Also, at the end of training, they were fixated less often. In other words, the neutral dimensions were at a double disadvantage in *classification* as well—they provided a only weak source of evidence which was largely ignored anyway. Models, like Baywatch and KRES, assume that information about features equally enter the network on every trial throughout training. But, the current eye fixations results indicate that these models are likely to overestimate the influence of neutral dimensions on final classification decisions because of continued input from those dimensions.

Knowledge selection in response to observed exemplars. The second question was the time course of knowledge effect on attention. It is appealing that the impact of knowledge will be greatest initially and then decrease with experience of category members, because prior knowledge is what learners *bring* to the learning task as compared to empirical observations that come later (e.g., Heit, 1995; Pazzani, 1991). However, the eye fixations results showed that subjects gradually allocated more eye fixations to the related dimensions with more experience with exemplars.

We interpret this result as arising from accumulation of semantic activations associated with the related features. That is, the knowledge associated with related features were not initially comprehensible, but repeated observation of sets of related features might have lead subjects to make sense of them via common theme (or via coherent relations among the related features). This process does not have to occur in a manner of "aha" experience because relations between related features could have been noticed in a piecemeal manner. Heit and Bott (2000) have labeled this "knowledge selection," and like us, showed gradual effect of knowledge on feature learning (but not on attention).

Baywatch and KRES assume that knowledge is in place from the start of training, rather than being constructed dynamically in response to observed category members. However, if knowledge were active from the start, these models predict initial attention preference for the related dimensions, a prediction we failed to confirm. Thus, because of the built-in knowledge representation, these models oversimplify the process by which knowledge is gradually activated in response to exemplar observations.

The (non)necessary role of error in attentional shift. The third question we asked is whether error is required to mediate attentional shift. Note that all current accounts of attention change during learning are based on error. For example, ALCOVE predicts gradual attention shift to dimensions that reduce error (Kruschke, 1992). But, contra this account, eye fixations continued to shift to related dimensions even after subjects learned to classify all items.

Thus, we conclude that error is not a necessary condition for attention shift in knowledge-based category learning.

We propose two explanations for this result. The first is what we have already mentioned, that is, the activation of semantic representations. In our experiment, merely observing exemplars that consist of related features might have been sufficient to further activate associated representations, that is, the tundra and desert themes. This theme discovery in the absence of error is consistent with the extensive literature documenting knowledge effect in unsupervised category learning. When prior knowledge is available, subjects can come up with theme-based (family resemblance) category construction.

The second reason that attention might shift without error is the desire for speed that our cognitive systems are trying to achieve—all else being equal, a faster response is more adaptive than a slower one. One way that latency can be decreased is by gathering less information in preparation of a decision, and of course, to maintain accuracy, less learned neutral dimensions were the first to go.

In supervised learning not involving prior knowledge, Rehder and Hoffman (2005a) also found that learners first discovered a one-dimensional rule and then, after a few errorless trials, they discontinued attending to other redundant dimensions. Moreover, Blair, Watson, & Meier (2009) found that learners continued to optimize attention in the absence of any feedback whatsoever. These results, along with the current one, altogether pose problems for all category learning models that tie attention to error-driven mechanisms (e.g., Kruschke, 1992).

Summary

Using eyetracking, we found that (a) knowledge directs attention to related dimensions and away from unrelated ones, (b) this effect did not emerge immediately but gradually emerged in response to observing category members, and (c) this effect grew even after the last error, that is, in the absence of error. Models of knowledge-based category learning will remain incomplete until they include attention mechanisms that explain these empirical results.

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant No. 0545298.

References

Blair, M., Watson, M. R., & Meier, K., M. (2009). Errors, efficiency, and the interplay between attention and category learning. *Manuscript submitted for publication*.

Bott, L., Hoffman, A., & Murphy, G. L. (2007). Blocking in category learning. *Journal of Experimental Psychology: General, 136*, 685-699.

Heit, E. (1995). Belief revision in models of category learning. In J. D. Moore & J. F. Lehman (Eds.), *Proceedings of the 17th Annual Conference of the Cognitive Science Society* (pp. 176-181). Mahwah, NJ: Erlbaum.

Heit, E., & Bott, L. (2000). Knowledge selection in category learning. In D. L. Medin (Ed.), *Psychology of learning and motivation* (pp. 163-199). San Diego: Academic Press.

Kaplan, A. S., & Murphy, G. L. (2000). Category learning with minimal prior knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26*, 829-846.

Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review, 99*, 22-44.

Kruschke, J. K. (1993). Three principles for models of category learning. In G. V. Nakamura, R. Taraban, & D. L. Medin (Eds.), *Categorization by humans and machines: The psychology of learning and motivation* (Vol. 29, pp. 57-90). San Diego: Academic Press.

Kruschke, J. K. (2001). Toward a unified model of attention in associative learning. *Journal of Mathematical Psychology*, *45*, 812-863.

Murphy, G. L., & Medin, D. L. (1985). The role of theories in conceptual coherence. *Psychological Review*, 92, 289-316.

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. (2002). *The big book of concepts*. Cambridge, MA: MIT Press.

Nosofsky, R. M. (1992). Exemplars, prototypes, and similarity rules. In A. F. Healy, S. M. Kosslyn, & R. M. Shiffrin (Eds.), *From learning processes to cognitive processes: Essays in honor of William K. Estes* (pp. 149-167). Hillsdale, NJ: Erlbaum.

Pazzani, M. (1991). The influence of prior knowledge on concept acquisition: Experimental and computational results. *Journal of Experimental Psychology: Learning, Memory & Cognition, 17*, 416-432.

Rehder, B., & Murphy, G. L. (2003). A Knowledge-Resonance (KRES) model of category learning. *Psychonomic Bulletin & Review, 10*, 759-784.

Rehder, B. & Hoffman, A.B. (2005a). Eyetracking and selective attention in category learning. *Cognitive Psychology, 51*, 1-41.

Rehder, B., & Hoffman, A. B. (2005b). Thirty-something categorization results explained: Selective attention, eyetracking, and models of category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31*, 811-829.

Rehder, B., Colner, R. M., & Hoffman, A. B. (2009). Feature inference learning and eyetracking. *Journal of Memory & Language, 60*, 394-419.

Shepard, R. N., Hovland, C. L., & Jenkins, H. M. (1961). Learning and memorization of classifications. *Psychological Monographs, 75*, (13, Whole No. 517).

Smith, E. E., Patalano, A. L., & Jonides, J. (1998). Alternative strategies of categorization. *Cognition, 65*, 167-196.