(Category) Learning by Doing: How Goal Directed Tasks Constrain Conceptual Acquisition

Seth Chin-Parker (chinparkers@denison.edu)
Department of Psychology, Denison University
Granville, OH 43023 USA

Abstract

The current study explores conceptual acquisition that occurs as the result of completing a task in a novel domain. The items encountered in the domain were complex in that there were multiple sources of information that might be used to organize conceptual knowledge related to the domain. I test the hypothesis that goal-directed interactions will constrain the acquisition of knowledge such that functional categories of the items, organized around goal-relevant features, are learned. Converging evidence from two measures provided strong support for the idea that participants organized their knowledge of the domain in terms of goal-relevant features, and the conceptual organization was able to support both the completion of the task and subsequent categorization tasks.

Keywords: category learning, goals, similarity.

Prior experience underlies intelligent behavior – people learn through interactions with the environment what behaviors lead to successful outcomes and what ones do not. An important component of this is recognizing categories of events and items among those experiences, a process that leads to the acquisition of conceptual knowledge, knowledge of those categories. That knowledge can be used to categorize, communicate, reason, and problem solve at later points. A fundamental question then is how coherent categories of items are identified so that the conceptual knowledge can be appropriately applied. Theories of categorization address what ties together items within a category and subsequently coheres the conceptual organization that reflects those categories, and most theories rely on a notion of similarity for at least a component of that cohesion (Hahn & Ramscar, 2001). Thus, the question shifts to how this similarity is determined.

There have been two basic approaches to answering this question (Malt, 1995). The first assumes that the environment constrains the similarity. Rosch et al. (1976) nicely capture this idea by positing that features of items in the world occur in reliable clusters and the conceptual system learns to recognize that structure. They go so far as to illustrate in one "experiment" (1976, Exp. 3) how overlaying tracings of the members of various basic level categories (e.g. cat, shoe, truck) results in greater perceptual overlap within-category than across categories. Similarly, others (e.g. Anderson, 1991) have stressed the role of the environment in determining conceptual structure. That view can be contrasted with one that places much of the emphasis on the individual to constrain the similarity. Murphy and Medin (1985) argued information from the environment has to be situated within the knowledge structures (e.g. inter-category relations and theories) that the individual brings to any interaction with the environment. In this way, the structure of conceptual knowledge is constructed as the individual interprets what is to be considered feature information and how those features relate to one another.

Although most researchers interested in concepts and categories stake out some middle ground in this debate, much of the work in human category learning assumes that the environment provides structure. This assumption has seemingly created a disconnect between work exploring more naturalistic concepts and the basic experimental work examining conceptual acquisition (Murphy, 2005). I identify two critical differences between basic experimental studies and more naturalistic ones and explore them in the current study. First, in most experimental work, the categories are well defined in terms of their features and structure. Second, participants interact with members of those categories with the goal of differentiating the items they encounter based on that structure. In more naturalistic studies, the presence of the categorical structure is less clear and people interact with the items not with the goal of classifying items, but with the goal of accomplishing some other task. I present a study that incorporates a more complex, arguably more naturalistic, structure and vary the interactions that participants have with the items. In this manner, I examine how goal-directed behaviors within a domain affect the structure of the conceptual knowledge acquired about items within that domain.

I begin this research with the assumption that the environment is not a source of simple, unambiguous information about the categories that exist. For instance, people are able to recognize and use information about the taxonomic categories of food items, e.g. breads and vegetables, but they also recognize and use goal-related categories, e.g. snack foods and breakfast foods, to guide inferences and determine appropriate groupings of foods (Ross & Murphy, 1999). Similarly, people can identify and use ad-hoc categories (Barsalou, 1991) to guide behaviors. A study by Medin, Lynch, Coley, and Atran (1997) illustrates how this complexity can be reflected in conceptual knowledge. The experimenters asked various tree experts to sort cards labeled with tree names into groups. Those experts concerned with research and teaching tended to create groups that were highly correlated with the biological taxonomy, but landscapers tended to create groups that reflected the way the trees would be incorporated into landscaping decisions (e.g. a shade tree versus a weed tree). These cross-classifications and the development of ad-hoc categories of items are problematic

for an account that posits that the environment alone provides structure to our conceptual knowledge (although see Anderson, 1991, for a rebuttal). Although we have evidence of the complexity of naturalistic categories, the structure of the categories used in basic experimental work does not reflect this complexity (Murphy, 2005). Most items that comprise the categories are defined by specific feature lists or simple visual features, and the relation of the features to the categories is also carefully controlled. This typically results in a structure with only one "correct" organization for the items. Instead of operating in a situation with multiple possible configurations, participants in experimental studies are placed into a situation that is much more constrained by the information available.

Within these experimental studies, the interactions that people have with the categories are also different from what occurs in more naturalistic situations. In a typical category learning study, an item is presented, the participant predicts the category membership of the item, and feedback is given on the classification judgment. This approach has produced a great deal of information about how people learn to classify items, but may not capture important aspects of how people learn about categories in more naturalistic situations (Ross, Chin-Parker, & Diaz, 2005). Numerous studies have shown that classification learning promotes a near exclusive focus on diagnostic information, the features that distinguish the categories (Chin-parker & Ross, 2004; Rehder & Hoffman, 2005), but it is not apparent whether other means of category learning share this restricted focus (e.g. Minda & Ross, 2004). Arguably, the interactions we have within more naturalistic contexts are more varied and richer than the classification decisions made in a typical experimental setting. Importantly, I note that these interactions occur not with the primary intention to learn about the categories but rather to accomplish some other goal. The importance of goal-directed interactions has been explored by a range of cognitive scientists (e.g. Ram & Leake, 1995), and goals are implicated to some extent in how we come to recognize structure in the environment (Love, 2005). For instance, the naturalistic studies mentioned prior (e.g. Medin, Lynch, Coley, & Atran, 1997) suggest goal-directed interactions give rise to conceptual organizations that are able to support those interactions.

So, it seems that we can begin to bridge the chasm between experimental and naturalistic study of concept acquisition by adopting more complex categorical structures and varying the goals of the participants as they interact with the items that comprise those categories. Recently, Jee and Wiley (2007) did just that. They had participants learn about creatures that could be distinguished in terms of their perceptual features, shown through simple line drawings, or the nutritional value and ability to avoid predators (information about these features was conveyed through a list of features located beneath the picture). In their study, participants initially organized items in terms of similarity of the simple perceptual features, but as they learned about the domain and interacted with items, they either learned to

identify the nutritional value or how the creature avoided predators, the goal-relevant information came to be important within the conceptual organization. Subsequent transfer tasks showed that a participant adopted a conceptual organization that reflected the information that was critical to their interactions within the domain, and the participants' similarity judgments were shaped by the presence of that information. Their study provides more clear evidence that the goal-directed interactions caused the shift in the conceptual structure.

In the current study, I examine category learning that occurs as the result of goal-directed interactions with items. Like Jee and Wiley (2007), I have a complex structure and participants interact with the items in accordance with different goals. In our study, the items are Flux Capacitor Boards, actual physical boards with various electrical components (non-functioning) affixed to them. As is described below, I created the boards so that there were two types of the boards that the participant would encounter during their initial task. However, only the classification participants were informed that these categories existed; the other participants were simply asked to complete their assigned task with the boards. Our primary hypothesis is that the conceptual organization adopted by the participants will be organized around the features of the boards that are relevant to the attainment of their goal.

As is described below, the *goal-relevant features* of the boards varied across the conditions. In one condition, the goal-relevant features are the configuration of specific components of the boards. In another condition, the goal-relevant features are relationships that exist between the components of the boards. For the classification condition, there were several possible sources of information that would be considered goal-relevant, or diagnostic. As noted in Jee and Wiley (2007), working towards a specific goal can often lead to information that is not goal-relevant to be left out of the conceptual organization. In this study, I expect that the two conditions with specific goal-relevant information will focus exclusively on that information, like classification learners in previous studies (e.g. Chin-Parker & Ross, 2004). Interestingly, since the classification condition will have multiple sources of information relevant to differentiating the categories, I predict that they will show a more general knowledge of the boards. I have no strong prediction as to whether the difference in the kind of goal-relevant information available to the two non-classification task conditions will affect the participants' acquisition of useful conceptual knowledge.

Experiment

Methods

Participants and Design Fifty-seven participants were randomly assigned to three experimental conditions: 18 participants were assigned to the *flexible condition*, 19 to the *solid condition*, and 20 to the *classification condition*. Two participants in the classification condition failed to show evidence of learning during their initial task, so their data

were removed from all analyses. All participants interacted with the same set of items during the initial task and completed the same two transfer tasks[1]. The presentation order of items during the initial and transfer tasks was randomized for each participant.

Materials and Procedure The primary materials for the study consisted of the Flux Capacitor boards and the connectors used to complete the boards. Each board had a series of nine terminal posts and various electrical components affixed to the board (see Figure 1). The posts were organized into three sets: One set in the upper, left-hand region of the board, one in the middle region, and one in the lower, right-hand region. The other components were placed around these posts according to the parameters described below. The boards were designed so that there were two types of boards that the participants encountered during the initial task, and variations of these two types of boards were created for the transfer tasks. During the initial task, participants in the flexible and solid conditions were given connectors that they placed onto the terminal posts to complete each board. The participants in the classification condition did not use connectors during their initial task

In the flexible condition, the connectors were made of wire and varied in terms of how they fit onto the terminal posts: The connector either fit over an open post or was inserted into a hole drilled into the "capped and drilled" post. As can be seen in Figure 1, each set of terminal posts in the Type A boards featured one post that has been capped and drilled and two posts that were open. In contrast, the Type B boards featured sets consisting of two capped and drilled posts and one open post. The configuration of the posts is considered to be the goal-relevant feature for the flexible condition because they constrain how the flexible connectors can be placed onto the board.

In the solid condition, the connectors were made of inflexible aluminum pieces, and the placement of these connectors was constrained by the presence of components situated near the terminal posts. For the Type A boards, the connectors had to go between components. For the Type B boards, the connectors had to go around the components. Thus, the relationship of the components to the posts is considered the goal-relevant feature for the solid condition because it constrained how the solid connectors could be placed onto the board.

The electrical components were unique to each board. However, each of the boards featured a perceptually salient *correlated component*. The correlated component for the Type A boards was a two-inch section of a computer memory module placed in the near left corner, and the correlated component for the Type B boards was a stack of silver clips with copper wire loops placed in the far right corner. These components were not implicated in how the connectors in either condition could be placed onto the board but were perfectly diagnostic of the two board types. During the initial task, participants in the solid and flexible conditions were asked to complete the boards by placing three of the six connectors onto the terminal posts. The participants in the classification condition were told that the boards were incomplete and before they could be completed they needed to be identified as "positive flux" or "negative flux" boards. The classification participants were instructed to learn how to identify the two types of boards. There were eight boards used in the initial task phase, half were Type A and half were Type B. Each participant encountered each board twice during this phase.

In the classification condition, a board was placed into the holder, and after the participant responded with either "positive flux" (correct for the Type A boards) or "negative flux" (correct for the Type B boards), the experimenter provided feedback about the classification and allowed the participant to study the board. In the solid and flexible conditions, the experimenter placed a board into the holder, and the participant determined which connectors to place

Figure 1: Example problem boards from Experiment 1

Notes: The boards on the left are Type A boards, and the boards on the right are Type B boards. The top images show boards with no operators present. The center images show boards completed with the flexible connectors. The bottom images show boards completed with the solid connectors.

---

[1] Participants also completed a sorting task following the same-different task. However, the classification condition was inadvertently given different instructions for the task, so we are unable to compare performance across the groups. The results of the task very closely tracked those of the same-different task.

onto the board. After each trial, the board was removed from sight and a new board was placed into the holder.

After the initial task, all participants completed the same two transfer tasks. The materials for the transfer tasks consisted of photographs of flux capacitor boards the participants had not encountered during the initial task. The boards in the images were designed so that they varied in terms of how they related to the Type A/Type B board distinction that had been present during the initial task. No feedback was given to participants as they completed the transfer tasks.

First, the participants completed the *same-different task.* During each trial, the participant was presented with images of two boards affixed to a piece of paper. She was asked to indicate whether she would consider the two boards pictured to be the same type or different types. Across the sixteen items in the same-different task, I balanced whether the boards matched or mismatched in terms of the goal-relevant features. Eight of the pairs of boards maintained the same structure as the initial tasks boards; four of those pairs matched and four mismatched. All participants regardless of condition should identify the matches as the same and the mismatches as different if they picked up on any of the sources of information that differentiated the Type A and Type B boards during the initial task. The other eight boards were designed so that the goal-relevant features from the solid and flexible conditions were placed into opposition. For instance, if the goal-relevant features for the flexible condition matched what had been seen on the Type A board, the goal-relevant features for the solid condition would match what had been seen on the Type B board. Four of these board pairs were designed so that flexible condition goal-relevant features matched while the solid condition goal-relevant features mismatched. The other four board pairs were designed so the flexible condition goal-relevant features mismatched while the solid condition goal-relevant features matched.

The *category goodness-rating task* was the final task. I balanced whether the Type A or Type B boards were rated first. The participant was first shown a target board, one of the boards solved during the initial task phase, and was told that the board was either an "X-12" (Type A) or "G-59" (Type B) board. She was asked to rate each subsequent board shown in terms of the category indicated by the target board on a scale from one ("excellent example of this board type") to nine ("not this type of board"); also anchored at three ("good example of this board type"), five ("ok example of this board type"), and seven ("poor example of this board type"). After the participant studied the target board for a minute, it was removed, and the items for the goodness-rating task were shown to the participant one at a time. There were five types of boards pictured in the stimuli for this task, and the participant rated two of each type for each of the categories. The *category consistent* boards were structurally identical to the target board. The *category inconsistent* boards were structured like the other type of board; so if the target board was a Type A board, the

category inconsistent board was a Type B board. The *correlation violation* boards were the same type of board as the target board, but the correlated feature was replaced by a small, perceptually dissimilar component. The *flexible violation* boards were of the same type as the target board, but were altered so the flexible connectors would not fit onto the posts. The *solid violation* boards were also of the same type as the target board, but they were altered so the solid connectors would not fit. Once the participant completed rating the ten boards for the first type, the target board for the second type was shown to the participant, and the task repeated for the second type.

Results

In the same-different task (Table 1), I found strong evidence that the participants in the flexible and solid conditions organized their knowledge of the domain in terms of the goal-relevant features. Across all items in the task, both the flexible condition, $M = 0.95$, $SD = 0.13$, $t(17) = 14.72$, $p < 0.001$, and the solid condition, $M = 0.87$, $SD = 0.18$, $t(18) = 8.64$, $p < 0.001$, were above chance performance in terms of assigning the pairs as the same or different in terms of the goal-relevant features for their conditions. The difference between the flexible and solid conditions was not significant, $t(35) = 1.57$, $p = 0.12$. A similar summarization of the results for the classification condition is not possible because there was no a priori prediction of how the classification participants would handle the items when the two goal-relevant features were placed in opposition. However, as can be seen in Table 1, when both of the goal-relevant features matched, they considered the boards as the same, and when both did not match, they considered the boards as different. When the goal-relevant features for the flexible and solid conditions were put into opposition (as in the "Flex + / Solid -" and "Flex - / Solid +" items), the participants in the classification condition did not show a preference for one source of information over the other as a group. Within the

Table 1: Proportion of Items (standard deviation) Identified as "the Same" in the Same-Different Task

	Relation of Boards in the Pair			
Condition	Flex + Solid +	Flex - Solid -	Flex + Solid -	Flex - Solid +
Flexible	0.94 (*0.24*)	0.01 (*0.06*)	0.90 (*0.26*)	0.04 (*0.18*)
Solid	0.86 (*0.21*)	0.11 (*0.23*)	0.15 (*0.29*)	0.86 (*0.21*)
Classification	0.83 (*0.33*)	0.19 (*0.24*)	0.46 (*0.39*)	0.36 (*0.36*)

Notes: For each item, the boards pictured either matched in terms of the goal-relevant features of the flexible (Flex +) or solid (Solid +) conditions or mismatched in terms of those features of the flexible (Flex -) or solid (Solid -) conditions.

classification condition, five participants had a pattern of response that indicated that they were using information about the goal-relevant features for the flexible condition, four participants appeared to be using information about the goal-relevant features for the solid condition, and nine participants had a pattern of responding that did not clearly indicate a preference for either source of information.

The category-goodness rating task provided a more specific indication of what information from the domain was being used by the participants in each condition. The data from the task (Figure 2) were analyzed using a series of ANOVAs. I report the results of five one-way ANOVAS that compared the ratings for each items type across the experimental conditions. I also include relevant within-condition comparisons where appropriate (full analyses are not included due to space restrictions).

There were no differences as to how participants in the three conditions rated the category consistent items, $F(2, 54) = 0.01$, $MSE = 0.02$, $p = 0.99$, but there were significant differences within all the other item types. Participants in all conditions rated the category inconsistent items as less good category members compared to all other items. Also, within the ratings for the category inconsistent items, there were some differences between the conditions, $F(2, 54) = 3.41$, $MSE = 1.27$, $p = 0.04$, primarily between the classification and flexible conditions. The ratings of the correlation violation items also varied by condition, $F(2, 54) = 3.72$, $MSE = 3.45$, $p = 0.03$. The classification condition rated these items as significantly worse category members than the category consistent items ($p < 0.01$) but the other two conditions did not. There were significant differences in the ratings of both the flexible violation items, $F(2, 54) = 21.62$,

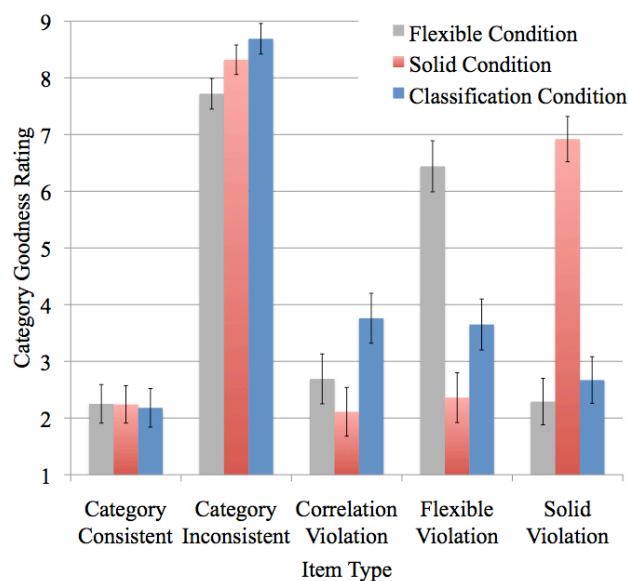Figure 2: Mean Category-Goodness Ratings by Condition and Item Type

Figure Note: The category-goodness rating scale ranged from one ("excellent example of this board type") to nine ("not this type of board").

$MSE = 3.69$, $p < 0.01$, and the solid violation items, $F(2, 54) = 40.51$, $MSE = 3.04$, $p < 0.01$. As predicted, the participants in the solid condition rated the solid violation items as significantly worse than the category consistent items ($p < 0.01$), but did not rate the flexible violation items differently than the category consistent items ($p = 0.54$). The participants in the flexible condition rated the flexible violation items as significantly worse category members than the category consistent items ($p < 0.01$), but not the solid violation items ($p = 0.83$). The classification condition did not rate the solid violation items as significantly worse than the category consistent items ($p = 0.13$), but did rate the flexible violation items as worse category members ($p = 0.02$).

Discussion

I return to the question regarding what constrains the similarity underlying conceptual organization. The results of the participants in the flexible and solid conditions clearly show that the goal-relevant features are central to their notion of similarity for the items and thus critical for the organization of categories of boards within the domain. They identify novel boards as the "same" when they match in terms of the goal-relevant features and "different" when those features do not match. They also show a pattern of category goodness ratings that indicates that violating those goal-relevant features makes the boards less good members of the category while violating other sources of information have little or no effect on those judgments. The participants in the classification condition seem to maintain a more diffuse attentional focus during the initial task. This is interesting given earlier studies that show a very narrow focus for classification learners. However, these results fit well together when we consider that the goal of classification learning is to predict the category membership of items. Typically only a subset of the information available within the experimental materials allows those judgments to be accurately made, so the classification learner attends most to that subset of information. In this study, multiple sources of diagnostic information existed, so the classification learners maintained a correspondingly wide attentional focus. It was the participants in the solid and flexible conditions that had a narrow focus in this experiment, and this was due to the fact that only a subset of the information available within the domain was goal-relevant for each condition.

In one sense the results of this study are not surprising – there are numerous models of learning that incorporate an attentional mechanism (as discussed in Kruschke, 2003) to account for shifts across the information available during learning. Although attention obviously plays a critical role in the learning, it is not a sufficient determinant of learning; we need to understand what drives the attention. I propose that attention is guided by comparisons between the boards as the participants interact with them in terms of the goal they have (how to place the connectors or classify the board), and those features that are relevant to the person's

goal are picked out during the comparison process (e.g. Medin, Goldstone, & Gentner, 1993). It is important to note that this approach helps to explain how "simple" perceptual features (e.g. the capped and drilled posts) and relational information (e.g. how components were positioned with regard to the posts) can both be picked up and used in grounding similarity.

One critical question is whether the participants in this study were *really* engaged in category learning, or to put it another way, did they *really* recognize categories of boards during the initial task? In the typical classification learning paradigm, this question is deflected because the participants explicitly know of the presence of the categories and their responses are made in response to those categories. However, as has been noted prior, there are questions as to whether even that *really* constitutes learning a category (Ross, Chin-Parker, & Diaz, 2005). For this study, I would argue that the participants acquire knowledge that is sufficient to support their task (identifying the "type" of board facilitates the placement of the connectors, and there was ample evidence of this facilitation occurring during the learning) and to guide later, more explicitly category-based tasks. It is at least the foundation of category learning.

The current study was not designed to address all facets of this process. For instance, additional study within this paradigm will be able to determine whether the participants came to adopt different representations of the features (e.g. Schyns, Goldstone, & Thibault, 1998) or whether they learned to ignore certain information (e.g. Denton & Kruschke, 2006) as they better discriminated the features during the learning. This paradigm provides a unique way to approach these types of questions and to situate the study of them within a larger framework intended to guide our understanding of the acquisition of conceptual knowledge.

Conducting this type of study within a more complex, and arguably more naturalistic, domain, we can begin to see the interaction between the individual and the environment that helps to shape the acquisition of conceptual knowledge. We can extend our understanding of the ways in which goals are implicated in category learning and how we might bridge the chasm that has separated naturalistic studies of concepts from more experimental studies.

References

Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review, 98*, 409-429.

Barsalou, L.W. (1991). Deriving categories to achieve goals. In G.H. Bower (Ed.), The psychology of learning and motivation: Advances in research and theory (Vol. 27, pp. 1-64). San Diego, CA: Academic Press.

Chin-Parker, S., & Ross, B. H. (2004). Diagnosticity and prototypicality in category learning: A comparison of inference learning and classification learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 30*, 216-226.

Denton, S. E., & Kruschke, J. K. (2006). Attention and salience in associative blocking. *Learning & Behavior, 34,* 285-304.

Hahn, U., & Ramscar, M. (Eds.). (2001). Similarity and categorization. New York : Oxford University Press.

Jee, B. & Wiley, J. (2007) How goals affect the organization and use of domain knowledge. *Memory & Cognition, 35*, 837-51.

Kruschke, J. K. (2003). Attention in learning. *Current Directions in Psychological Science, 12*, 171-175.

Love, B. C. (2005). Environment and goals jointly direct category acquisition. *Current Directions in Psychological Science, 14,* 829-835.

Malt, B. C. (1995) Category Coherence in Cross-Cultural Perspective. *Cognitive Psychology, 29,* 85-148.

Medin, D. L., Goldstone, R. L., & Gentner, D. (1993). Respects for similarity. *Psychological Review, 100,* 254-278.

Medin, D. L., Lynch, E. B., Coley, J. D., & Atran, S. (1997). Categorization and reasoning among tree experts: Do roads lead to Rome? *Cognitive Psychology, 32*, 49-96.

Minda, J. P., & Ross, B. H. (2004). Learning categories by making predictions: An investigation of indirect category learning. *Memory & Cognition, 32*, 1355-1368.

Murphy, G. L. (2005). The study of concepts inside and outside the laboratory: Medin versus Medin. In W. Ahn, R. L. Goldstone, B. C. Love, A. B. Markman, & P. Wolff (Eds.), Categorization inside and outside the lab: Festschrift in honor of Douglas L. Medin. Washington, DC: APA.

Murphy, G. L., & Medin, D. L. (1985) The role of theories in conceptual coherence. *Psychological Review, 95*, 289-316.

Ram, A., & Leake, D. B. (Eds.). (1995). Goal-driven learning. Cambridge; MIT Press.

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

Rosch, E., Mervis, C. B., Gray, W., Johnson, D., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology, 8,* 382-439.

Ross, B. H., Chin-Parker, S., & Diaz, M. (2005). Beyond classification learning: A broader view of category learning and category use. In W. Ahn, R. L. Goldstone, B. C. Love, A. B. Markman, & P. Wolff (Eds.), Categorization inside and outside the lab: Festschrift in honor of Douglas L. Medin. Washington, DC: APA.

Ross, B. H., & Murphy, G. L. (1999). Food for thought: Cross-classification and category organization in a complex real-world domain. *Cognitive Psychology, 38*, 495-553.

Schyns, P. G., Goldstone, R. L., & Thibaut, J-P (1998). Development of features in object concepts. *Behavioral and Brain Sciences, 21*, 1-54.