Learning mode and comparison in relational category learning

John D. Patterson (jpatter4@binghamton.edu)
Kenneth J. Kurtz (kkurtz@binghamton.edu)
Department of Psychology, 4400 Vestal Parkway East
Binghamton, NY 13905 USA

Abstract

An important goal in the study of higher-order cognition is to understand how relational categories are acquired and applied. Previous work has explored the potential of within-category comparison opportunities to promote relational category learning and transfer. This follows from predictions of structure mapping theory (Gentner, 1983, 2003) that alignment leads to highlighting and abstraction of common relational structure. However, a straightforward merging of traditional classification learning with comparison (i.e., trials presenting two same-category items) has not been effective. We explore the hypothesis that classification and comparison have an unforeseen incompatibility. In a 3x2 between-subjects design we tested three presentation conditions (unconstrained item pairs, category-matched items pairs, single items) in two supervised category learning modes: classification and observation. The major finding is an interaction driven by highly accurate categorization for the observational learners with same-category pairs. The introduction of the observational mode yielded the predicted, but elusive result of an advantage for within-category pairs over twice as many single-item trials. We conclude that within-category comparison can be an effective means to promote relational category learning and discuss the apparent impediment of the guess-and-correct cycle.

Keywords: relational categories; structural alignment; comparison; classification learning; transfer; observational learning mode

Introduction

Categorization and comparison are two core mechanisms underlying human learning, comprehension, and knowledge use. Within the study of categorization, the bulk of the research attention has been devoted to object/entity categories – categories whose members belong based on sharing a set of intrinsic features. Though the learning and generalization of entity categories has been studied using different tasks, such as through inference of missing features (Markman & Ross, 2003) and category construction (Ahn & Medin, 1992), the traditional classification learning paradigm has remained the most prevalently used (Murphy, 2003; Ross, Chin-Parker, & Diaz, 2005). In its most common form, the traditional classification learning paradigm operates as follows: a single stimulus is presented, the participant is asked to classify the item into one of two category options, a response is selected, and corrective feedback is given. The traditional, classification learning paradigm has yielded substantial knowledge and has offered a testing ground for formal models of categorization (e.g., ALCOVE, Kruschke, 1992; DIVA, Kurtz, 2007; SUSTAIN, Love, Medin, & Gureckis, 2004).

However, not all categories lend themselves well to traditional accounts of category learning. Though a sizable amount of knowledge can be captured through a feature-based understanding of the world, features alone do not seem to capture the richness of what we know – the ways in which objects and attributes relate to one another reflect a critical facet of the concepts we hold. In the categorization literature, an increasing emphasis has been placed on relational categories (Gentner & Kurtz, 2005; Markman & Stilwell, 2001) that are based on a common (perhaps rule-like) relational structure as opposed to a set of intrinsic features (see Corral & Jones, 2014; Goldwater, Markman, & Stilwell, 2011; Higgins & Ross, 2011; Kurtz, Boukrina, & Gentner, 2013; Patterson & Kurtz, 2014; Smith & Gentner, 2014; Weitnauer, Carvalho, Goldstone, & Ritter, 2014). As an example, take the relational noun *bridge* – something that connects two other things. A member of the category *bridge* might occupy the form of a concrete structure connecting two landmasses. Alternatively, a *bridge* might take the form of an ambassador connecting the geopolitical ideas of two countries. In terms of features, these two members of the category *bridge* are greatly disparate; a bridge does not have much in common with a diplomat. Nonetheless, both *bridges* are category members insofar as they relate to two other things in the same way. This qualitative difference between entity and relational categories translates to differences on the quantitative level as well, with relational categories exhibiting slower acquisition in children (Gentner, 2005). These differences expose an empirical need for the study of relational categories in order to further the understanding of human categorization.

A pressing topic in the study of relational categories is how they are learned. Prior investigation has demonstrated benefits to the acquisition and transfer of relational categories through comparison (Kurtz, Boukrina, & Gentner, 2013; Patterson & Kurtz, 2014). The observed comparison benefits can be understood through the process of structural alignment (Gentner, 1983, 2003, 2010; Gentner & Markman, 1997). According to the structural alignment view, deep, relational similarity that exists between two cases is rendered salient by aligning their relational predicates during comparison, allowing for shared relational structure to be abstracted into a portable knowledge structure. It is predicted from this view that comparison advantages should be great when same-category items are compared, relative to single item learning and comparison using contrasting categories whose relational structures are not alignable. Although same-category comparison during classification learning has been shown to confer benefits on near and far transfer of category knowledge (relative to twice as many trials of single item learning), the comparison advantage has only been found with the inclusion of some different-category pairs (Kurtz, Boukrina, & Gentner, 2013). A fifty-fifty mix of same and different-category pairs led to comparison outperforming single item learning. One important characteristic of this format is that the learner cannot assume both items belong to the same category and is therefore encouraged to consider each of the items and their category assignment relative to one another. While a comparison advantage was found, it was not clear what specific aspect(s) of this methodology provided the causal power. To date, we know of no successful demonstration of a pure, same-category comparison advantage over single item learning – one successful attempt required twice as many stimulus exposures as the single item control (Kurtz & Gentner, 1998). In preliminary work to the current study, we attempted to boost the invitation to compare during same-category comparison trials. However, neither the use of similarity ratings nor the drawing of correspondence lines between compared items yielded differences between pure, same-category comparison and single item learning. In these cases it seemed clear that the comparison engine was effectively engaged but, perplexingly, no advantage accrued.

An alternate account of these observed shortcomings is that the task acted in opposition to the benefits of same-category comparison. Drawing on the machine learning and attribute-based categorization literatures, a continuum can be found between two different learning modes: discriminative and generative learning (Levering & Kurtz, 2015; Ng & Jordan, 2001). Discriminative learning is characterized by learning the probability of a category given a set of features; the focus is on learning a minimalist way to predict a category given aspects of the stimulus. By contrast, a generative mode emphasizes learning the probability of a set of features given a category; in other words, the focus is on learning what stimulus aspects are common to a given category, resulting in a more positively defined, holistic representation. Consistent with this generative/discriminative distinction, previous work has shown the discriminative guess-and-correct cycle of classification to result in reduced holistic category knowledge compared to a more generative learning mode (Levering & Kurtz, 2015). Accordingly, the less holistic category representation encouraged by classification might be at odds with making productive comparisons. From a more general standpoint, a conflict may exist between performing classification and getting the most out of comparison – such that successfully coordinating and integrating the two components is not possible.

An alternative to the traditional classification-learning paradigm is supervised observational learning. On each trial, items are simply presented with their correct category labels. While the two modes generally lead to similar performance outcomes (Estes, 1994; see also, Ashby, Maddox, & Bohil, 2002; Edmunds, Milton, & Wills, 2014), observational learning has been shown to result in richer category knowledge (Levering & Kurtz, 2015). Using unidimensional rule plus family resemblance categories, Levering and Kurtz (2015) found observational learners showed enriched knowledge of internal category structure (relative to classification learners), demonstrating enhanced ability to infer values on the partially diagnostic features when provided only the category. Further, typicality ratings revealed greater sensitivity to changes on partially diagnostic features for observational learners compared to classification learners. Applied to relational categories, the more holistic consideration encouraged by observational learning could provide benefits to comparison and relational discovery.

Observational learning presents a viable task alternative to circumvent potential impediments associated with pure, same-category comparison learning under classification. The task allows the learner to jointly consider the co-presented examples as members of a category without the distractions of the guess-and-correct cycle. It is expected that, through unhindered structural alignment, the greatest benefit at test and far transfer will be conferred to same-category comparison in the observational mode, relative to mixed comparison (having half as many same-category comparison opportunities) or single item learning.

Method

The purpose of the experiment was to assess the impact of learning mode on the effectiveness of different kinds of comparison opportunities. To accomplish this, learning mode and presentation condition were varied orthogonally.

Participants

184 undergraduates from Binghamton University participated for partial course credit.

Materials

The training and testing phase stimuli consisted of 36 unique, Stonehenge-like arrangements of rocks – examples can be seen in Figure 1. Rocks varied in their size, shape, and color. As in our previous studies, the stimuli comprised three relational categories (category labels in brackets): *monotonicity* [Besod] – defined by a monotonic decrease in height of the arrangement from left to right, *support* [Makif] – characterized by the presence of a rock being supported by two other rocks, forming a sort of bridge, and *symmetry* [Tolar] – captured by the presence of two same color rocks of similar size and shape, one stacked atop the other. Each arrangement belonged to only one of the three categories. Of the 36 stimuli, a subset of 24 was utilized as the training set (eight per category) and 12 were reserved for use at test (four per category). The subsets matched those used in Kurtz, Boukrina, and Gentner (2013) and subsets were held constant across participants. For comparison conditions, training stimuli were presented in pairs. Pairs were randomly generated for each participant according to the condition – all same-category pairs (Same_ conditions) or a fifty-fifty blend of same- and different-category pairs (Mix_ conditions).

To assess far transfer of category knowledge, a set of 15 mobile-like stimuli (colorful, geometric objects connected with vertical lines, as if hanging down from a platform; see Figure 1) was used. Each mobile conformed to one of the three relational categories from training, five mobiles per category. Compared to the training and testing stimuli, the mobiles were dissimilar in their surface characteristics (in color and shape of objects) and the orientation of the category-defining core in each item was reflected over the X-axis.

Procedure

In a between-subjects design, participants were randomly assigned to one of six conditions. Four conditions employed comparison learning: same-category classification (SameClass, n = 30), mixed-category classification (MixClass, n = 31), same-category observational (SameObs, n = 31), and mixed-category observational (MixObs, n = 32). Two conditions served as single item controls: single item classification (SingClass, n = 31) and single item observational (SingObs, n = 29). All participants received an archeological cover story and were given the following instructions: "Your overall goal is to figure out what makes a given rock arrangement belong to one of the three types: Besods, Makifs, or Tolars. You will be tested on your knowledge of each type later." The following instructions were given to comparison conditions (and were stripped of dual-item and comparison language in the single item conditions): "On each learning trial, you will see two rock arrangements. [*Obs*: You will be shown the correct type for each arrangement to help you learn, *Class*: Try to figure out the correct type for each arrangement. Use the mouse to select your response. A box will appear around the arrangement that you should respond to. You will be given feedback at the end of each trial to help you learn]. At first you will not understand what makes them belong to a type, but before long you should become quite good at recognizing the different types. Remember that there are three different styles for arranging the rocks into configurations. Looking at the two arrangements together can help you learn these types. Try your best to gain mastery of the names of each type and what makes an arrangement belong to those types. Learn as much as you can before the test!"

Comparison Conditions – Training

Training consisted of two cycles of 12 paired stimulus trials, totaling 48 stimulus exposures. At the beginning of each trial, two laterally offset stimuli were presented and remained visible until the trial was complete. In the classification conditions, a box appeared that randomly queried one of the arrangements. Participants were asked for the category of the queried item. They selected a response using the mouse and were then queried about the other item. Following both responses, participants were shown simultaneous feedback for each item indicating: (1) whether or not their response was correct, (2) the correct category of the item (in green), and (3) if incorrect, the category they responded with (in red). In the observational conditions, the correct category labels were provided with the presented items and remained on screen for the duration of the trial. When the participant finished studying an item pair they continued to the next trial with a mouse click. Participants in both classification and observational conditions had as much time to engage each trial as they wished.

Single Item Conditions – Training

Training consisted of two cycles of 24 randomized, single item trials, totaling 48 stimulus exposures. A single stimulus was presented at the start of each trial and remained visible until the trial was complete. Classification and observational conditions closely followed their comparison counterparts. In the classification condition, participants were asked for the category of the item. Following their response, they were presented feedback identical in nature to the comparison classification conditions. In the observational condition, participants were presented with a single labeled item. As in the comparison conditions, single item conditions were permitted as much time as desired on each trial.

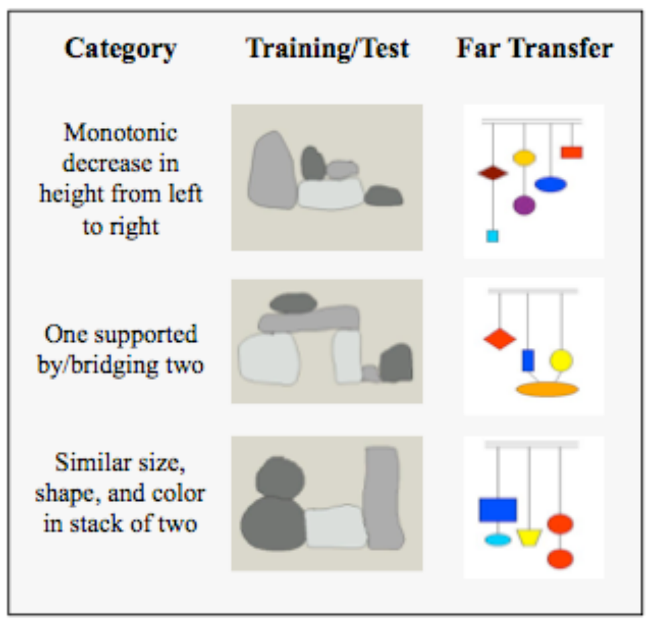

Figure 1: Sample stimuli for each category in each phase.

Assessment

Following training, all conditions performed an identical assessment sequence. The sequence consisted of, first, a within-domain test and, second, a far transfer assessment.

The within-domain test randomly presented the 24 "old" rock arrangements from training and 12 new arrangements.
After the within-domain test, the 15 mobile stimuli were presented in random order for the far transfer phase. Both the test and transfer trials employed an endorsement format: on each trial a single item was presented, the participant was asked if the item belonged to a given category, and participants gave a yes/no response. This measure of categorization performance is similar to, but distinct from both the classification and observational learning tasks. The endorsement task minimizes any transfer appropriate processing advantages (Morris, Bransford, & Franks, 1977) that might result from a perfect task match between training and testing phases. As the primary interest was in how well knowledge could be extended from the learning phase, old test items were each presented once, while the new test and far transfer items were each presented twice – once each with accurate and inaccurate category labels.

Results

Given the absence of training accuracy data for observational conditions, the analyses for training data are omitted here. While our predictions primarily concern the extension of knowledge to new examples, we begin by considering performance on old test items.

Test – Old Items

In the absence of complete training data, old test data can give an estimation of late learning phase performance. The data were subjected to a two-way analysis of variance with two levels of task (classification and observational learning) and three levels of presentation format (same-category pairs, mixed-category pairs, and single item). The ANOVA revealed a significant effect of task, $F(1, 178) = 7.71$, $p = .006$, showing that observation learners ($M = 0.84$, $SD = 0.15$) were more accurate in their endorsement decisions than were classification learners ($M = 0.77$, $SD = 0.17$). No main effect was found for presentation format, $F(2, 178) = 0.72$, $p = .49$, indicating that, collapsed across task, type of comparison did not have an effect. Consistent with our predictions however, a significant interaction showed that task differentially impacted the effectiveness of the type of comparison opportunity, $F(2, 178) = 3.77$, $p = .025$. The interaction was marked by a significant difference between observational ($M = 0.88$, $SD = 0.13$) and classification ($M = 0.74$ $SD = 0.19$) learning modes for same-category pairs [$t(59.36) = -3.42$, p = .001, corrected for unequal variances], but only a marginal difference between observational ($M = 0.85$, $SD = 0.13$) and classification ($M = 0.79$, $SD = 0.14$) learning modes for mixed-category pairs, $t(61) = -1.91$, p = .06. Task did not did not have an effect on single item learning (SingClass, $M = 0.79$, $SD = 0.18$; SingObs, $M = 0.78$, $SD = 0.18$; $p > .1$).

Single Item Control As a reminder, one of the primary goals of the experiment was to explore whether pure, same-category comparison could lead to an advantage over single item learning. As predicted, same-category comparison in the observational mode was found to provide a significant benefit over its task-matched, single-item control using a t-test comparison – the only comparison condition to do so, $t(53.10) = 2.45$, $p = .02$, corrected for unequal variances.

Test – New Items

Looking at the data for never-before-seen, within-domain items (see Figure 1, Figure 2), a 3x2 ANOVA denoted only a significant interaction, $F(2, 178) = 4.85$, $p = .009$. Follow up analyses indicated that a same-category observational performance advantage over its classification counterpart drove the interaction: while task led to significant differences in endorsement accuracy for same-category comparison conditions (SameClass, $M = 0.69$, $SD = 0.17$; SameObs, $M = 0.83$, $SD = 0.13$; $t(53.90) = -3.70$, $p = .001$, corrected for unequal variances), task did not lead to reliable differences between mixed-category comparison conditions or single item conditions. These results emphasize the power of pure same-category comparison, but only under the appropriate task circumstances.

Single Item Control As seen with the old-item test data, same-category learning in the observational mode was found to produce the predicted advantage over its task-matched, single-item control on new test items [$t(48.09) = 3.22$, $p = .002$, corrected for unequal variances]. By contrast, no other comparison condition was able to show an advantage over single item learning.

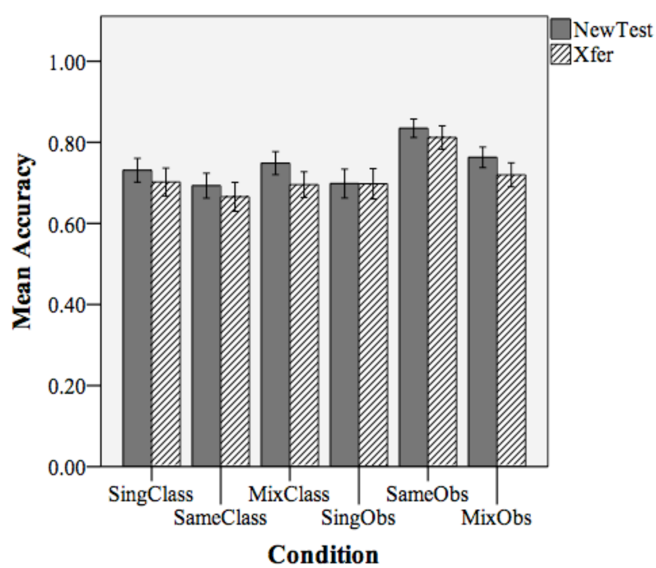

Figure 2: New item test and far transfer endorsement accuracy by condition. Error bars show +/- 1 SE.

Far Transfer

Of critical interest was the impact different learning conditions had on the ability to transfer category knowledge. To assess this, the far transfer accuracy data were subjected to a 3x2 ANOVA that indicated a main effect of task: learners in observational learners (*M* = 0.77, *SD* = 0.17) showed enhanced transfer relative to classification learners (*M* = 0.68, *SD* = 0.19), *F*(1, 178) = 4.22, *p* = .04. Consistent with both old and new test items, no main effect of presentation format was found on far transfer items. However, a marginally significant interaction between task and presentation format was found, *F*(2, 178) = 2.93, *p* = .056. Consistent with our predictions, the interaction showed that levels of task resulted in larger disparities between same-category comparison conditions (SameClass: *M* = 0.67, *SD* = 0.20; SameObs: *M* = 0.81, *SD* = 0.16; *t*(59) = -3.20, *p* = .002) than between mixed pair comparison conditions (MixClass: *M* = 0.70, *SD* = 0.18; MixObs: *M* = 0.72, *SD* = 0.17; *t*(61) = -0.56, *p* = .58). No differences were seen between single item learning conditions (SingClass: *M* = 0.70, *SD* = 0.20; SingObs: *M* = 0.70, *SD* = 0.19; *t*(58) = -0.09, *p* = .93).

Single Item Control Comparing same-category comparison in the observational mode against its single item, task-controlled counterpart yielded the predicted advantage for comparison at far transfer, *t*(58) = 2.43, *p* = .02. No other comparison condition exhibited reliable differences over single item learning.

Discussion

The empirical goal of the present study was to further the understanding of how relational categories are best learned. The specific questions being asked were: (1) how does learning mode affect the acquisition of relational categories? and (2) how does learning mode influence the effectiveness of different types of comparison opportunities? The results show clearly that the observational mode has a positive influence on learning, increasing endorsement accuracy on within-domain members and enhancing far transfer to members across domain. While the type of comparison opportunity did not exhibit a direct impact on performance, exceptional performance in the same-category observational group drove an interaction at within-domain test and pushed an interaction near significance at far transfer. The interaction underscores that the type of learning task plays an important role in the effectiveness of certain kinds of comparison opportunities (same-category pairs), but not for others (mixed-category pairs). Further, same-category comparison in the observational mode was the only condition to display an advantage over single item learning – and did so across all testing phases.

These results are compelling for a number of reasons. First, they represent the first time that observational learning has been shown to outperform feedback learning on a test of category membership knowledge. Though the effectiveness of observation as a learning vehicle has been sparsely explored, previous work has shown observational learning to be either equivalent or disadvantaged relative to classification when category membership is the target of assessment (Ashby, Maddox, & Bohil, 2002; Edmunds, Milton, & Wills, 2015; Estes, 1994; Levering & Kurtz, 2015). It should be noted that most research employing observational learning has been conducted using feature-based categories with single item presentation. Taken together, this poses the possibility that the type of category (feature-based or relational) may interact with task and presentation format; this is a topic for further research. Second, the findings clearly echo that learning mode can substantially impact acquired category knowledge (Markman & Ross, 2003). This highlights the need for future categorization research to study phenomena using a broader palette of learning methods. Third, this study demonstrates a pure, same-category comparison advantage over single item learning for the first time. This key finding fits nicely into the theoretical framework developed in the study of analogy (Gentner, 1983; Markman & Gentner, 1997). The success of same-category comparison under observation, relative to classification, suggests that classification is disruptive to fruitful comparison. Future work will seek to further specify and elaborate on this finding.

The excellent level of mean performance in the same-category observational condition is unprecedented in the study of relational category learning. Accordingly, understanding the basis for this success is paramount. A number of causal factors are worth exploring. One speculation is that observational learning encourages greater engagement than classification (despite being a less active task: there is no responding). Unlike classification, observational learning with same-category pairs does not involve the guess-and-correct cycle. As such, classification may promote discriminative goals that interfere with making the most of comparison opportunities. Classification learners may be more inclined to look for diagnostic features and less attuned to relational structure. Also, classification learners may be more focused on the performance factor of getting correct answers as opposed to the more global goal of category mastery.

A Further Speculation

One speculative factor that may serve to benefit engagement in the same-category observational case is *symbolic juxtaposition* (Gentner, 2010). By applying the same label to the presented items, it represents an invitation through language to compare and abstract commonalities that exist between them. Getting this invitation at the beginning of the trial (as opposed to the feedback period at the end of classification trials) might have emphasized comparison as the focus of the task and led to more engaged and effective comparisons.

References

Ahn, W., & Medin, D. L. (1992). A two-stage model of category construction. *Cognitive Science*, 16(1), 81-121.

Ashby, F. G., Maddox, W. T., & Bohil, C. J. (2002). Observational versus feedback training in rule-based and information integration category learning. *Memory &*

Cognition, 30, 666–677.

Corral, D. & Jones, M. (2014). The effects of relational structure on analogical learning. *Cognition*, 132 (3), 280-300.

Edmunds, C.E.R., Milton, F., & Wills, A.J. (2015). Feedback can be superior to observational training for both rule-based and information-integration category structures. *Quarterly Journal of Experimental Psychology*, 68, 1-20.

Estes,W.K.(1994). *Classification and cognition*. Oxford: Oxford University Press.

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7, 155–170.

Gentner, D. (2003). Why we're so smart. In D. Gentner & S. Goldin-Meadow (Eds.), *Language in mind: Advances in the study of language and thought* (pp. 195–235). Cambridge, MA: MIT Press.

Gentner, D. (2005). The development of relational category knowledge. In L. Gershkoff-Stowe & D. H. Rakison (Eds.), *Building object categories in developmental time*. Hillsdale, NJ: Erlbaum.

Gentner, D. (2010). Bootstrapping the mind: Analogical processes and symbol systems. *Cognitive Science*, 34, 752–775.

Gentner, D., & Kurtz, K. J. (2005). Relational categories. In W. K. Ahn, R. L. Goldstone, B. C. Love, A. B. Markman, & P. W. Wolff *Categorization inside and outside the laboratory: Essays in honor of Douglas L. Medin* (pp. 151–175). Washington, DC: American Psychological Association.

Gentner, D., & Markman, A. B. (1997). Structure mapping in analogy and similarity. *American Psychologist*, 52(1), 45-56.

Goldwater, M. B., Markman, A. B., & Stilwell, C. H. (2011). The empirical case for role-governed categories. *Cognition*, 118, 359–376.

Higgins, E. J., & Ross, B. H. (2011). Comparisons in category learning: How best to compare for what. In L. Carlson, C. Holscher, & T. Shipley (Eds.), *Proceedings of the 33rd Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.

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

Kurtz, K. J. (2007). The divergent autoencoder (DIVA) model of category learning. *Psychonomic Bulletin & Review*, 14, 560–576.

Kurtz, K. J., Boukrina, O., & Gentner, D. (2013). Comparison promotes learning and transfer of relational categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 39(4), 1303–1310.

Kurtz, K. J., & Gentner, D. (1998). Category learning and comparison in the evolution of similarity structure. *In Proceedings of the Twentieth Annual Conference of the Cognitive Science Society* (p. 1236). Mahwah, NJ: Erlbaum.

Levering, K. R., & Kurtz, K. J. (2015). Observation versus classification in supervised category learning. *Memory & Cognition*, 43(2), 266-282.

Love, B. C., Medin, D. L., & Gureckis, T. M. (2004). SUSTAIN: A Network Model of Category Learning. *Psychological Review*, 111(2), 309-332.

Markman, A. B., & Ross, B. H. (2003). Category use and category learning. *Psychological Bulletin*, 129(4), 592-613.

Markman, A. B., & Stilwell, C. H. (2001). Role-governed categories. *Journal of Experimental & Theoretical Artificial Intelligence*, 13, 329– 358.

Morris, C. D., Bransford, J. D., & Franks, J. J. (1977). Levels of processing versus transfer appropriate processing. *Journal of Verbal Learning and Verbal Behavior*, 16, 519–533.

Murphy, G. L. (2003). The downside of categories. *Trends In Cognitive Sciences*, 7(12), 513-514.

Ng, A.Y., & Jordan, M. (2001). On discriminative vs. generative classifiers: A comparison of logistic regression and naive Bayes. In T. G. Dietterich, S. Becker, and Z. Ghahramani, editors, *Advances in Neural Information Processing Systems*, 14, 841–848.

Patterson, J.D., & Kurtz, K.J. (2014). Engaging the comparison engine: Implications for relational category learning and transfer. Poster presented at the 36th Annual Conference of the Cognitive Science. Quebec City, Quebec.

Ross, B. H., Chin-Parker, S., & Diaz, M. (2005). Beyond Classification Learning. In W. 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. 197-213). Washington, DC, US: American Psychological Association.

Smith, L. A., & Gentner, D. (2014). The role of difference-detection in learning contrastive categories. *Proceedings of the Thirty-Sixth Annual Conference of the Cognitive Science Society*. (pp. 2088-2093). Quebec City, Quebec: Cognitive Science Society.

Weitnauer, E., Carvalho, P.F., Goldstone, R.L., & Ritter, H. (2014). Similarity-based Ordering of Instances for Efficient Concept Learning. In P. Bello, M. Guarini, M. McShane, & B. Scassellati (Eds.), *Proceedings of the 36th Annual Conference of the Cognitive Science Society* (pp. 1760-1765). Austin, TX: Cognitive Science Society.