

Improving Category Learning Through the Use of Context Items: Compare or Contrast?

Janet K. Andrews (andrewsj@vassar.edu)

Program in Cognitive Science, 124 Raymond Avenue, Box 146
Poughkeepsie, NY 12604 USA

Kenneth R. Livingston (livingst@vassar.edu)

Program in Cognitive Science, 124 Raymond Avenue, Box 479
Poughkeepsie, NY 12604 USA

Kenneth J. Kurtz (kkurtz@binghamton.edu)

Department of Psychology, P.O. Box 6000
Binghamton, NY 13902 USA

Abstract

In four experiments, participants were trained to classify organism-like visual stimuli into three categories. On each training trial, the target item was presented with two other items varying in category membership. Learning was superior when each item in the triples was a member of a different category, though the strength of this effect depended on the nature of the categories being learned. In addition, there was an overall benefit of explicit prior knowledge of the triples structure. These results suggest, contrary to exemplar models, that abstraction processes do occur during category learning and, contrary to prototype models, that information about category contrasts is more helpful than information about category commonalities.

Keywords: Categorization; classification; concepts; learning.

Introduction

While there is a vast body of research on concepts and classification, and many issues have been investigated in depth, there has been relatively little exploration of the role of the context in category learning. One important kind of context is the set of comparison items within which a to-be-learned item is embedded. Comparison of examples has been shown to influence category learning in research that includes work on explicit reminders to previous cases (Ross, Perkins, & Tenpenny, 1990; Spalding & Ross, 1994), array-based presentation of entire item sets (Regehr & Brooks, 1995), and pairwise comparison of examples of relationally-defined categories (Kurtz & Gentner, 1998; Kurtz & Boukrina, 2004). The finding by Kurtz and Boukrina (2004) of some improvement in classification accuracy when training with pairs versus individual items points to at least one kind of positive context effect on category learning. Comparison of examples has also been shown to promote knowledge change in numerous domains that elicit structured representations (e.g., Gick & Holyoak, 1983; Loewenstein & Gentner, 2001).

We investigate the role of stimulus juxtaposition in classification learning by presenting target items in the context of coordinated triads. This "triples" paradigm allows us to study variously structured triads that invite the learner to

conduct: 1) comparative evaluation of within-category examples; 2) contrastive evaluation of between-category examples; or 3) both. Other variables that are also manipulated are the explicitness of information given about the nature of the item triples and the nature of the category definitions. Effects of context may only occur when learners are made aware of the nature of the context, or with certain kinds of category structure. An important further goal is to continue efforts to advance a more naturalistic basis for the study of category learning by using multiply-instantiated feature values (Markman & Maddox, 2003) of concrete perceptual stimuli, and three-way rather than binary classification (Homa & Vosburgh, 1976; Kurtz & Boukrina, 2004). An argument can also be made that in naturally-occurring category learning and categorization situations, it is not uncommon for instances of the same and/or contrast categories to be available at one time.

The project is novel in the following ways: 1) The ability of learners to leverage valuable information contextually embedded in coordinated sets of training instances has, to our knowledge, never been studied. 2) The use of flat-featured perceptual stimuli extends the study of the effects of inter-item comparison on learning beyond the domain of structural alignment theory (Gentner & Markman, 1997). 3) The influence of item juxtaposition on learning may help differentiate theoretical accounts of categorization. If exemplar models (e.g., Medin & Schaffer, 1978; Nosofsky & Palmeri, 1997) are correct in de-emphasizing the role of summary representations of category concepts, then manipulating presentation context should have little impact on category learning. If abstraction of commonalities among labeled category members drives learning, consistent with prototype models (e.g., Hampton, 1995; Rosch, 1978), then contexts which afford within-category comparison should be beneficial. If discrimination of differences between members of competing classes is a critical factor in category formation, as might be the case for hypothesis testing and decision boundary models of perceptual categorization (e.g., Nosofsky & Palmeri, 1998) then the opportunity to assess contrasts between members of different categories should be a facilitative factor.

Experiment 1

Six conditions are included, varying in the kind of information presented during training. In five of these conditions, a triple of items was presented on each trial. In the aAA condition, both context items match the category of the target item ‘a’; in the aAB condition, one context item matches the category of the target item and one does not; in the aBB condition, both context items are from the same category which differs from the target item’s category; in the aBC condition, all three items are from different categories; and in the aXX condition, the context items are random with respect to category. A single item control (SIC) condition was structured in the same way as the aXX condition, but each item was presented one at a time.

One possible learning strategy is to locate common features between items known to belong to the same category and perform abstraction, in which case the aAA condition should be advantaged. Another potential strategy is to identify contrasts between items belonging to different categories, in which case the aBC condition is most beneficial. The aBB condition offers weaker forms of both abstraction and differentiation on every trial, while the aAB condition is least informative since the learner doesn’t know for certain which context item matches the target item’s category and which does not.

Method

Participants A total of sixty-nine Vassar College undergraduates participated in partial fulfillment of an introductory psychology research requirement.

Stimuli The stimuli consisted of organism-like patterns created in Adobe Photoshop that varied systematically along three dimensions: body-aspect ratio, flagella length, and stripe width (see Figure 1). Body-aspect ratio consisted of a series of ovals with the same area but different dimensions, progressing gradually from more elongated to more round. Flagella consisted of a group of wavy lines that increased in size. Stripes mimicked the shading properties of a sine wave gradient and increased in width. For each dimension, eight values were selected such that adjacent values were just clearly distinguishable based on informal psycho-physical testing.

Three categories were defined using the higher or lower four dimension values as shown in Table 1.

Table 1: Category definitions for Experiment 1.

Category	Body-aspect ratio	Flagella length	Stripe width
Gex	More elongated	Shorter	Wider
Kij	More elongated	Longer	Narrower
Zof	Rounder	Longer	Wider

The structure on the left of the stimulus remained constant across all stimuli. Note that all three dimensions are necessary to learn these categories.

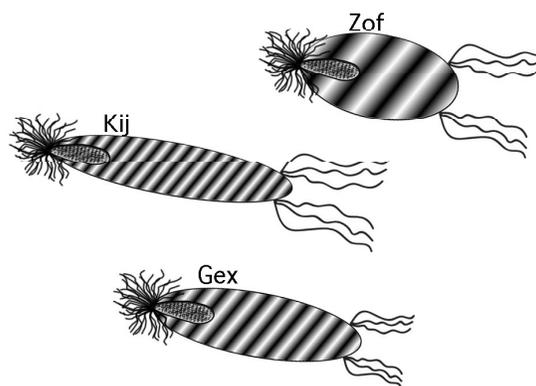


Figure 1: Sample stimuli for Experiment 1.

Procedure Each participant was randomly assigned to one of the six conditions, each of which consisted of a training phase and a test phase, both conducted using SuperLabPro software on Macintosh iMac computers. For all conditions, the training phase consisted of an instruction screen describing the nature of the condition followed by 144 trials. For example, for participants in the aAA condition, the screen would indicate that the three objects shown on each trial would always be from the same category. For conditions aAA, aAB, aBB, aBC, and aXX, on each trial, a triple of items configured as shown in Figure 1 would appear for 3 seconds followed by the identical screen with a red arrow indicating which of the three items was the target. The participant would press a key to classify the target as a Gex, Kij, or Zof and would immediately receive a feedback screen indicating if the response was correct or not, and if not, giving the correct classification. The SIC condition was exactly like the aXX condition except that each item was shown individually rather than in triples.

Of the 64 possible stimuli in each category, 48 were used an equal number of times, once as a target and twice as a nontarget (context) item. No item appeared in a triple with itself and no two items appeared more than once in a triple. The location of the target item in the triple varied randomly, as did the order of the triples. For all conditions, in the test phase, participants classified all possible 192 items presented one at a time in random order without feedback. At the conclusion of the test phase, participants were asked what they paid attention to and what strategies they used, if any, to learn the categories. Space limitations preclude discussion of those responses here.

Results and Discussion

To examine learning during the training phase, the 144 trials were divided into six blocks of 24 trials. A 6 (condition) by 6 (block) analysis of variance with repeated measures on block on proportion correct yielded significant effects of both condition ($F(5,63) = 3.324$, $MSE = .174$, $p = .01$) and block ($F(5,315) = 43.868$, $MSE = .014$, $p < .0001$). The interaction

between condition and blocks was nearly significant ($F(5,315) = 1.508$, $MSE = .014$, $p = .0592$). Performance improved substantially over blocks though less so for the aAB condition. Post hoc analysis (Fisher's PLSD) indicated that condition aBC produced significantly more accurate performance than all others, condition aAB produced significantly less accurate performance than all others, and condition aAA was significantly less accurate than conditions aBB and aEC as well as aBC. All groups perform well above chance, which is .33.

A 6 (condition) by 3 (category) analysis of variance with repeated measures on category on overall proportion correct in the training phase produced significant effects of both condition ($F(5,63) = 3.392$, $MSE = .086$, $p = .0089$) and category ($F(2,126) = 18.223$, $MSE = .012$, $p < .0001$). The Kij category produced more accurate classification, suggesting that stripe width is the most salient dimension.

Table 2: Accuracy in training phase of Experiment 1.

Condition	Mean	SD
aAA	.66	.22
aAB	.55	.24
aBB	.74	.19
aBC	.83	.16
aXX	.69	.18
SIC	.73	.16

For accuracy in the test phase, a 6 (condition) by 3 (category) by 2 (old/new) analysis of variance with repeated measures on the last two variables yielded significant effects of category ($F(2,120) = 30.713$, $MSE = .047$, $p < .0001$) and old/new ($F(1,60) = 24.858$, $MSE = .006$, $p < .0001$). The Kij category remains most accurately classified and old items ($M = .729$, $SD = .235$) are classified slightly more accurately than new items ($M = .690$, $SD = .245$), which are classified well above chance, indicating that participants were able to apply their knowledge of the categories to novel instances. There was also a significant interaction between category and old/new ($F(2,120) = 5.458$, $MSE = .005$, $p = .0054$), with new Zof instances being least accurately classified. The effect of condition was not significant ($p = .1886$), but the pattern of means was similar to that of the training phase, with BC being most accurate ($M = .793$, $SD = .161$) and AB being least accurate ($M = .603$, $SD = .292$).

Experiment 1 suggests that the nature of the triples presented during category learning does affect classification accuracy, at least for the duration of the training. A concern was the possibility that the information given at the start of training about the nature of each condition might not have been retained. To disambiguate this variable, two additional experiments were conducted. In Experiment 2, no information at all was given about the special nature of each condition, while in Experiment 3, participants were told the information and then given a series of pre-training trials to test their retention of the information.

Experiment 2

Method

Participants A total of fifty-five Vassar College undergraduates participated in partial fulfillment of an introductory psychology research requirement.

Stimuli and Procedure These were identical to those of Experiment 1 except that the initial information screen did not identify the nature of the triples to be presented, hence was identical for all conditions. Because the SIC condition of Experiment 1 presented no information that could be removed, this condition was not run again. Thus Experiment 2 included only conditions aAA, aAB, aBB, aBC, and aXX.

Results and Discussion

The analyses performed for Experiment 1 were performed on the accuracy data of Experiment 2 and a very similar pattern of results emerged. For the training data, there were significant effects of both condition ($F(4,50) = 4.194$, $MSE = .123$, $p = .0053$) and block ($F(5,250) = 43.126$, $MSE = .013$, $p < .0001$). Unlike Experiment 1, there was no interaction between condition and blocks since the AB condition improved as much as the others, perhaps because participants in this condition were less aware of its added complexity. Post hoc analysis (Fisher's PLSD) indicated that condition aBC again produced significantly more accurate performance than all others and in addition performance in condition aAA was significantly worse than in condition aAB.

A 6 (condition) by 3 (category) analysis of variance with repeated measures on category on overall proportion correct in the training phase produced significant effects of both condition ($F(4,50) = 4.282$, $MSE = .059$, $p = .0047$) and category ($F(2,100) = 12.632$, $MSE = .015$, $p < .0001$), as in Experiment 1.

Table 3: Accuracy in training phase of Experiment 2.

Condition	Mean	SD
aAA	.61	.22
aAB	.67	.16
aBB	.66	.19
aBC	.83	.14
aXX	.64	.17

The test phase analysis of variance on proportion correct yielded the same pattern of results as in Experiment 1, with significant effects of category ($F(2,100) = 38.394$, $MSE = .042$, $p < .0001$), old/new ($F(1,50) = 14.010$, $MSE = .006$, $p = .0005$), and category by old/new interaction ($F(2,100) = 15.580$, $MSE = .005$, $p < .0001$). Once again the effect of condition was not significant ($p = .2443$), but the pattern of means was similar to that of the training phase, with condition aBC being most accurate ($M = .817$, $SD = .194$). However, here condition aAA is the least accurate ($M = .668$, $SD = .257$).

The similarity between the results of Experiments 1 and 2 suggests either that prior knowledge of the triples structure has no effect on category learning accuracy, or that the information given to participants in Experiment 1 was insufficient to demonstrate an effect of prior knowledge. The purpose of Experiment 3 was to determine which of those is the case by strengthening the nature of the information given to participants.

Experiment 3

Method

Participants A total of fifty-one Vassar College undergraduates participated in partial fulfillment of an introductory psychology research requirement.

Stimuli and Procedure These were identical to those of Experiment 2 except for the following: Prior to the training phase, participants were informed of the structure of the triples in their condition and given ten trials presenting incomplete triples composed of simple geometric shapes. On each trial, participants indicated what shape the missing item could or had to be and received feedback on their responses. Because there is no systematic structure to the triples in the aXX condition that participants can be pre-trained on, this condition was not included. Thus Experiment 3 had only conditions aAA, aAB, aBB, and aBC

Results and Discussion

Pre-training performance differed considerably across conditions; subsequent analyses were carried out only on data from participants performing better than chance (15/15 in aAA, 6/12 in aAB, 8/12 in aBB, and 10/12 in aBC).

The training phase data for Experiment 3 produced a significant effect of blocks ($F(5,185) = 29.109$, $MSE = .014$, $p < .0001$) but not condition ($p = .1514$), though the pattern of means was fairly close to that of Experiment 1 (see Table 2). Similarly, for overall proportion correct in the training phase, there was only an effect of category ($F(2,74) = 10.913$, $MSE = .013$, $p < .0001$).

Table 4: Accuracy in training phase of Experiment 3.

Condition	Mean	SD
aAA	.78	.17
aAB	.69	.14
aBB	.68	.21
aBC	.82	.16

The test phase data for the participants in Experiment 3 who performed well on the pre-training showed the familiar effects of category ($F(2,74) = 16.345$, $MSE = .032$, $p < .0001$) and old/new ($F(1,37) = 11.407$, $MSE = .004$, $p = .0017$), though no interaction of the two. There is again no effect of condition ($p = .5899$).

While the reduced number of participants in several conditions lessens the power of the analysis, the results of Experiment 3 suggest that the effect of condition – in particular, the advantage of the aBC condition seen clearly in the training phases of Experiments 1 and 2 – is essentially removed when structured triples are presented and participants understand that structure. Comparing the results of Experiment 3 to those of Experiment 2, it appears that condition aAA performance was improved by the pre-training manipulation.

Experiment 4

Experiments 1-3 taken together suggest that under some circumstances, category learning is improved by the presentation of items in triads consisting of a target and two other items, where the three are all from different categories. Although the significance of this effect was not maintained during the testing phase or when all participants fully understood the structure of the various types of triples, it did occur despite exactly the same items being presented in the different conditions, and only the configuration of presentation varying.

It is interesting that the superior condition was aBC, which highlights differences between categories. This might seem somewhat surprising in light of the emphasis in the category learning literature on similarities and abstraction within categories. The purpose of Experiment 4 was to test whether aBC remains advantageous using a very different category structure than that used in Experiments 1-3. In particular, we expected that using several category-irrelevant dimensions of variation and one, relatively subtle category-relevant dimension would favor comparative evaluation of within-category examples, i.e., the aAA condition, over the aBC condition. Experiment 4 included only these two types of triples, along with a prior information manipulation.

Method

Participants A total of sixty-six Vassar College undergraduates participated in partial fulfillment of an introductory psychology research requirement.

Stimuli The stimuli were similar to those used in Experiments 1-3 but category membership depended on only one dimension, flagella length, which took on the same eight values used previously and one additional, shorter value. Gexes had the shortest three values, Kijes the middle three, and Zofs the longest three. The eight values of body-aspect ratio and stripe width used previously were also used but varied randomly, with all values occurring in all categories. Two new dimensions each took on four values and also varied randomly: the size of the structure that appeared on the left of all stimuli in Experiments 1-3, and the width of the oval's outline. Examples are shown in Figure 2.

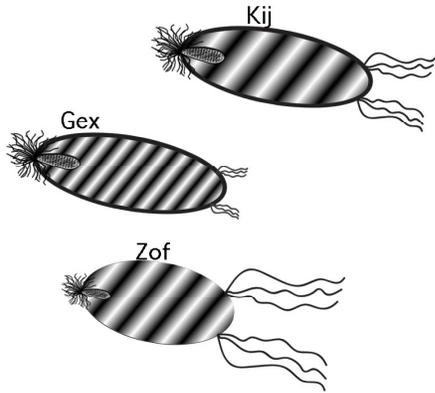


Figure 2: Sample stimuli for Experiment 4.

Procedure Only conditions aAA and aBC were used but each occurred both with no information given to participants about the triples structure (comparable to the instructions used in Experiment 2) and with information and pre-training on the triples structure (comparable to the procedure used before the training phase in Experiment 3).

Since there are many more possible instances of these categories, a set of 48 instances of each category was chosen randomly with the constraint that equal numbers of every dimensional value were used and no two features were correlated with each other, within or across categories. These 48 instances from each category were used in the training phase, which was otherwise identical procedurally to the training phase of Experiments 2 and 3. For the test phase, and to avoid chance imbalances in the sampling from the three categories, an additional 16 items were randomly chosen from the center of each category by using only the middle flagella value for each category and the middle two values of each irrelevant dimension. These additional items were used along with the training phase items in the test phase, which was procedurally identical to that used in Experiments 1-3.

In addition to the aAA and aBC conditions with no prior information, and the aAA and aBC conditions with pre-training on the triples structure, a single item control condition (SIC) was also included for purposes of comparison in learning and classification performance.

Results and Discussion

These categories were more difficult to learn than those used in Experiments 1-3. Pre-training performance was good for both the aAA and aBC conditions, so no data were excluded from the analyses. Analysis of training phase accuracy without the SIC condition allowed for condition and prior information to be incorporated as independent variables. This analysis showed significant effects of condition ($F(1,48) = 19.030$, $MSE = .225$, $p < .0001$), information ($F(1,48) = 4.350$, $MSE = .225$, $p = .0423$), and block ($F(5,240) = 14.426$, $MSE = .020$, $p < .0001$), and a significant condition by block interaction ($F(5,240) = 7.811$, $MSE = .020$, $p < .0001$) because performance in the aAA conditions barely improved over the training.

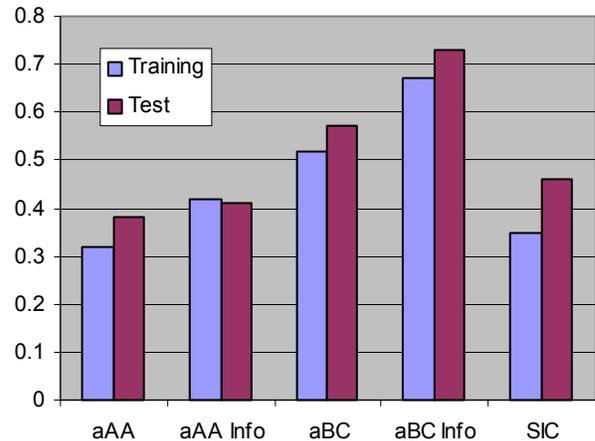


Figure 3: Mean proportion of correct classifications in training and test phases of Experiment 4.

Similar analysis of overall proportion correct in the training phase yielded significant effects of condition ($F(1,49) = 17.288$, $MSE = .023$, $p = .0001$) and information ($F(1,49) = 5.175$, $MSE = .023$, $p = .0273$). Post hoc analysis (Fisher's PLSD) with the SIC condition included indicated that all five means differ significantly except that neither aAA condition differs significantly from SIC. Mean performance in several conditions remains at or near chance.

Analysis of test phase accuracy produced significant effects of condition ($F(1,49) = 23.499$, $MSE = .218$, $p < .0001$) and category ($F(2,98) = 3.471$, $MSE = .038$, $p = .0350$), with Kij responses being slightly more accurate. The condition by information by old/new interaction was also significant ($F(1,49) = 9.335$, $MSE = .023$, $p = .0036$) due to old (but not new) items being classified more accurately in the aAA information condition than the aAA no information condition. The main effect of information did not reach significance ($F(1,49) = 3.033$, $MSE = .218$, $p = .0879$). Post hoc analysis (Fisher's PLSD) with the SIC condition included indicated that all five means differ significantly except that the aAA information condition did not differ significantly from either aAA no information or SIC.

Experiment 4 thus produced a pattern of results similar to that shown in Experiments 1-3, but much more strongly in that the aBC condition was superior for both training and test phases compared to both aAA and SIC conditions, and whether participants were explicitly pre-trained to ensure prior knowledge of the triples structure or not. Unlike the previous experiments, prior information about the triples structure did not remove differences between the aAA and aBC conditions, though it did tend to improve performance in both.

General Discussion

The experiments reported here demonstrate that success in classification learning can be affected by the nature of the items presented in coordinated triads during training, and suggest that learners are able to benefit most from such triads when each item is a member of a different category. For the

set of categories used in Experiments 1-3, this effect was only significant in the absence of clear prior knowledge of the structure of the triads, and only during training. For the set of categories tested in Experiment 4, triads containing members of all three categories led to significantly better classification accuracy than triads containing members of the same category, and this was true for both training and subsequent testing with old and new items. It was also true whether participants were explicitly informed and pre-trained on the structure of the triads prior to learning or not, though prior information improved performance during training. It is interesting that the random variation and relatively subtle category distinctions used in Experiment 4 produced an even stronger advantage for the different-category triads over the same-category triads.

These findings suggest that the potency of inter-item comparison during learning that has been demonstrated in numerous domains involving structured representations can potentially be extended to classification of flat-featured perceptual stimuli, and may have implications for the relationship between structural alignment principles and categorization (Lassaline & Murphy, 1998). It is also important to consider how our findings relate to available evidence concerning sequential item comparisons that may be occurring during traditional single-item classification training.

These results are of additional interest because they are difficult to reconcile with exemplar or prototype models of category learning. In exemplar models, learning depends on the set of labeled examples that are presented, which was kept constant over all conditions of our experiments, and on selective attention to diagnostic features, and all features were equally diagnostic in our first three experiments. Thus it is not clear how these models can explain the observed aBC advantage. Prototype models allow for some information beyond labeled instances and attention weights to be derived during learning, information to which the inter-item comparison afforded by our triads could potentially be relevant. However, such models would predict that the aAA condition would be most beneficial since it promotes the abstraction of common elements. We can therefore conclude that there is an abstractive process beyond what exemplar models allow, but that, surprisingly, it is better supported (at least in the case of our design and materials) by the opportunity to easily pick out category-differentiating information rather than within-category commonalities. Exploration of other kinds of models will be needed to clarify alternative approaches to category learning that will more successfully capture this pattern of findings.

Acknowledgments

Our thanks to the Undergraduate Research Summer Institute at Vassar College, and to Peter Alfaro, Ariana Brill, Jason Chen, Jessica Cicchino, Elizabeth Gonsiorowski, and Emma Myers for their assistance with these experiments.

References

- Gentner, D., & Markman, A. (1997). Structure mapping in analogy and similarity. *American Psychologist*, *52*, 45-56.
- Gick, M., & Holyoak, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology*, *15*, 1-38.
- Hampton, J. A. (1995). Testing the prototype theory of concepts. *Journal of Memory and Language*, *34*, 686-708.
- Homa, D., & Vosburgh, R. (1976). Category breadth and the abstraction of prototypical information. *Journal of Experimental Psychology: Human Learning and Memory*, *2*(3), 322-330.
- Kurtz, K. J., & Boukrina, O. (2004). Learning relational categories by comparison of paired examples. *Proceedings of the Twenty-Sixth Annual Conference of the Cognitive Science Society*.
- Kurtz, K. J., & Gentner, D. (1998). Category learning and comparison in the evolution of similarity structure. *Proceedings of the Twentieth Annual Conference of the Cognitive Science Society* (p. 1236).
- Lassaline, M. E., & Murphy, G. L. (1998). Alignment and category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *24*(1), 144-160.
- Loewenstein, J., & Gentner, D. (2001). Spatial mapping in preschoolers: Close comparisons facilitate far mappings. *Journal of Cognition and Development*, *2*(2), 189-219.
- Markman, A. B., & Maddox, W. T. (2003). Classification of exemplars with single- and multiple-feature manifestations: The effects of relevant dimension variation and category structure. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *29*(1), 107-117.
- Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review*, *85*, 207-238.
- Nosofsky, R. M., & Palmeri, T. J. (1997). An exemplar-based random walk model of speeded classification. *Psychological Review*, *104*, 266-300.
- Nosofsky, R. M., & Palmeri, T. J. (1998). A rule-plus-exception model for classifying objects in continuous-dimension spaces. *Psychonomic Bulletin and Review*, *5*(3), 345-369.
- Regehr, G., & Brooks, L. R. (1995). Category organization in free classification: The organizing effect of an array of stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*(2), 347-363.
- Rosch, E. (1978). Principles of categorization. In E. Rosch & B. B. Lloyd (Eds.), *Cognition & categorization* (pp. 27-48). Hillsdale, NJ: Erlbaum.
- Ross, B. H., Perkins, S. J., & Tenpenny, P. L. (1990). Reminding-based category learning. *Cognitive Psychology*, *22*, 460-492.
- Spalding, T. L., & Ross, B. H. (1994). Comparison-based learning: Effects of comparing instances during category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*(6), 1251-1263.