Learning Within-Category Attribute Correlations in a One-Attribute Visual Search Classification Paradigm

Gyslain Giguère (giguere.gyslain@courrier.uqam.ca)
Université du Québec à Montréal, Département de psychologie, c/o LEINA
C.P. 8888, Succ. CV, Montréal, Qc, H3C 3P8, CANADA

Guy L. Lacroix (guy_lacroix@carleton.ca)
Glen Howell (ghowell@connect.carleton.ca)
Carleton University, Department of Psychology,
1125 Colonel By Dr., Ottawa, On, K1S 5B6, CANADA

Serge Larochelle (serge.larochelle@umontreal.ca)
Université de Montréal, Département de psychologie,
C.P. 6128, Succ. CV, Montréal, Qc, H3C 3J7, CANADA

Abstract

In this study, participants categorized stimuli in a one-attribute rule visual search classification paradigm. The stimuli were six-shape displays that included a rule attribute and five diagnostic attributes. In Experiment 1, attribute values were changed at transfer. Slower RTs were obtained when attribute values from conflicting categories were used. In Experiment 2, the rule attribute (and up to two other attributes) were removed at transfer. The results showed that several attributes (color, texture, and size) of varying diagnosticity were used to correctly classify the stimuli. These experiments provide evidence that within-category attribute correlations can be learned in a classification task without intentional or inference learning instructions.

Keywords: Category learning; Classification learning; Within-category correlations; Incidental learning; Explicit learning.

Introduction

Classification Learning

Through the years, results from the classification learning paradigm have consistently led researchers to assert that participants typically adopt a hypothesis-testing strategy which involves trying simple, one-attribute rules (Bruner, Goodnow, & Austin, 1966; Nosofsky, Palmeri, & McKinley, 1994). Nevertheless, when the categories to be learned do not allow the use of a one-attribute rule, participants must pay attention to many attributes in order to successfully categorize stimuli, thus longitudinally creating exemplar memories. The resulting representations will generally be incomplete, however, as the memories only encompass the attributes involved in making categorical judgments (i.e., the diagnostic attributes), instead of complete stimuli (Chin-Parker & Ross, 2004).

This conclusion is further supported by a paradigm that has studied the impact of exemplar similarity on the application of a classification rule (Lacroix, Giguère, & Larochelle, 2005; Regehr & Brooks, 1993). In these studies, participants were given a complex disjunctive rule to sort creature-like stimuli belonging to two categories. Each creature had five idiosyncratically implemented attributes, three of which were specified in the rule. Results showed that transfer stimuli that were highly similar to training stimuli in one category, while belonging to the opposite category by virtue of the rule, elicited higher error rates and longer response times than other transfer stimuli. Regehr and Brooks suggested that the effect stemmed from both rule and non-rule attributes. However, Lacroix et al. (2005) showed that the two non-rule attributes could be modified without consequence for the categorization results, and that the effects stemmed from learning the association between the perceptual aspects of the attributes specified in the rule and category membership. Hence, once more, these results lead to the conclusion that, generally, it is the attention that is paid to the stimulus attributes that determines what is learned during the categorization process, and that participants are not naturally inclined to learn relationships among attributes (i.e. within-category correlations) unless their attention is focused on them.

Within-category attribute correlation learning

A few studies, however, have sought to show that within-category correlations can be learned. For instance, Anderson and Fincham (1996, Experiments 2 & 3) asked their participants to categorize schematic flowers composed of two features (petal and sepal) that varied on two dimensions (length and shading). Then, in a prediction phase, the participants had to determine one dimension from the remaining three. The results showed that they had developed sensitivity to the within-category correlations. Nevertheless, the composition of the stimuli only required participants to focus their attention on two features in order to process the four dimensions (as length and shading were integrated into a single feature). Moreover, the participants had to classify the stimuli and graphically reproduce them after each trial during the categorization phase. Anderson and Fincham acknowledged that without this additional demand, participants did not pay attention to both features. Hence, it is unclear whether participants were learning the within-category correlations via the categorization task or whether they were simply memorizing them.

Using a similar paradigm, Thomas (1998) asked his participants to classify exemplars belonging to two categories built from two correlated continuous dimensions (circle size and orientation of a radial line). Then, one of the two dimensions was removed and participants were asked to predict the appearance of the missing dimension. Again, this study showed that participants had been able to learn the within-category correlations. Once more, however, the attention paid to these integrated stimuli allowed participants to process both dimensions simultaneously.

In conclusion, the studies by Anderson and Fincham (1996) and Thomas (1998) do not allow us to ascertain whether within-attribute correlations can be learned during a categorization task when the stimuli are comprised of a large number of perceptually-separable features.

Inference Learning

In fact, the only task that has successfully shown consistent learning of within-category attribute correlations when multiple perceptually-separable features are used is inference learning[1] (Anderson, Ross, & Chin-Parker, 2002; Chin-Parker & Ross, 2002; 2004), a task in which participants are asked to infer an exemplar's missing attribute from its remaining attributes and category label.

Chin-Parker and Ross (2004) provide a strong example of this line of research. Their study compared participants' acquisition of categorical knowledge in classification and inference learning tasks. During learning, participants were presented creature-like stimuli composed of five binary attributes belonging to two family resemblance categories. At transfer, participants were asked to provide typicality judgments for new labeled instances. The category members presented at transfer differed in the number of diagnostic and prototypical attributes, a diagnostic attribute being one that distinguished between the two categories whereas a prototypical attribute was one that was shared within a category. The authors found that participants who performed an inference learning task relied more on prototypical than on diagnostic attributes to perform the typicality rating task. The reverse result was obtained with participants in the classification learning condition. Thus, Chin-Parker and Ross concluded that participants performing an inference task acquire knowledge of within-category attribute correlations (i.e., family resemblances) and to a lesser extent of between-category attribute correlations (i.e., diagnostic attributes), whereas participants performing an induction task almost exclusively acquire knowledge of between-category attribute correlations.

Nevertheless, in Chin-Parker and Ross's (2004) experiments participants had to focus on the within-attribute correlations to successfully predict attributes during inference learning. Participants who did the classification learning task could successfully categorize without doing so. Hence, the design of the stimuli, in combination with the task demands, forced the participants to infer the relationships among attributes in the inference learning task.

The One-Attribute Rule Classification Task

In order to determine if correlations among attributes could be learned without explicitly or tacitly focusing the participants' attention on them, Giguère, Lacroix, and Larochelle (2007) conducted a one-attribute rule classification experiment. The authors suspected that Lacroix et al. (2005) had been unable to find evidence that participants learned non-rule attributes because the rules used were complex and burdensome on participants' attention. Also, non-rule attributes were non-diagnostic and they were located at the periphery of the stimuli.

Participants in Giguère et al.'s (2007) study were assigned a single attribute rule to classify a set of fictional creatures. Each creature consisted of seven binary attributes: one perfectly diagnostic rule attribute (Rule), one attribute that was always correlated (PC) with the Rule attribute, and five family resemblance attributes (FR) that were partly correlated (0.80) with the rule. The authors varied the position of the rule and non-rule attributes within the creatures. There were global attributes spanning the entire stimulus (e.g., color), and local attributes confined to a part of the stimulus (e.g., number of legs). The saliency of the rule and non-rule attributes was also varied. The high salience attributes were color and body type and the low salience attributes were texture and number of legs.

During the training phase, participants classified the fictional creatures based on the assigned one-attribute rule. During the first test phase, participants classified a set of new fictional creatures for which the Rule attribute was removed. During the second test phase, participants classified a set of new fictional creatures in which both the Rule and PC attributes were removed. The authors expected that participants would classify the transfer items with above chance accuracy if they acquired knowledge of non-rule attributes. Transfer results showed that participants were significantly more accurate at classifying the fictional creatures in the first test phase, when the Rule attribute was global (the creature's texture) and the PC attribute was global and salient (the creature's color). The authors concluded that the participants' knowledge of the PC attributes was learned incidentally. This study provided evidence that within-category correlations could be incidentally learned when participants are engaged in rule-driven classification.

[1] We must note that Lassaline and Murphy's (1996) participants also acquired this type of categorical knowledge with perceptually-separable features, but the task demands of their experiment explicitly required participants to learn within-category attribute correlations.

The One-Attribute Rule Visual Search Categorization Paradigm

Although novel, Giguère et al.'s (2007) finding that the participants can incidentally learn non-rule attributes was limited to a single condition, one involving a global and salient attribute that was perfectly correlated with the Rule attribute. The other transfer test items, which did not contain the Rule or PC attributes, were classified at chance level in all conditions. Thus, participants were unable to incidentally learn the within-category correlations of FR attributes. The participants' potential for incidental learning of non-rule attributes may have been restricted because the rule attributes appeared in the same location within the stimulus on every trial. Consequently, participants may not even have looked at local, non-rule attributes.

The goal of the present study was to test whether participants can incidentally learn within-category attribute correlations in a paradigm that forces the participants to visually scan the stimulus, namely a visual search paradigm. The experimental design was similar to Giguère et al.'s (2007), but differed in four important ways. First, each stimulus was composed of six shapes displayed separately instead of single fictional creatures. Second, the six shapes were randomly displayed around the center of the screen so that participants had to search the display for the shape specified in the one-attribute classification rule (See Smith, Redford, Gent, & Washburn (2005) for a recent example of a visual search categorization task). We reasoned that the visual search aspect of our design could lead to more learning of within-category correlations than in Giguère et al.'s study. Third, we created six experimental conditions by varying the physical characteristics that served as the Rule, PC, and FR attributes. Finally, because higher amounts of training have been shown to increase the effect of non-rule attribute learning in rule-driven classification (Lacroix et al., 2005), the participants in this study were subjected to four times as many training trials (640 trials) than Giguère et al.'s participants (160 trials).

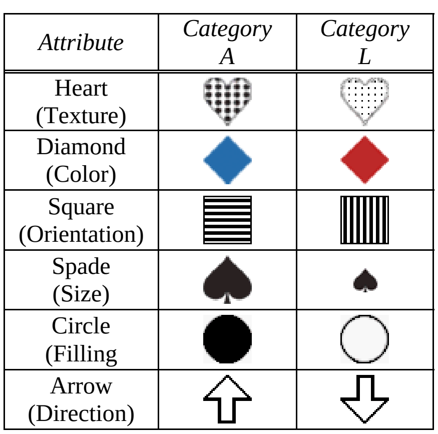

Figure 1: Visual implementation of attribute values associated with each category (Experiments 1 and 2)

Experiment 1

The rationale behind this first experiment is as follows: if participants acquire knowledge about the within-category correlations through extensive training, then violating some of these correlations at test should slow down categorization response times (RT). RTs constitute a very sensitive, on-line, indirect measure of the incidental learning of a correlational structure. Hence, if the search component of our categorization task allows participants to incidentally learn the association between the rule attribute and other attributes, then transfer items that break these associations should generate longer RTs.

Method

Participants 60 undergraduate students from the Université de Montréal participated in the experiment. They received 8$ as compensation for their time.

Material The eight training phase stimuli were hexagonal six-shape displays, where each shape represented a binary attribute (Figure 1): *heart* (big spots vs. small spots), *diamond* (blue vs. red), *square* (horizontal vs. vertical lines), *spade* (large vs. small), *circle* (full vs. hollow), or *arrow* (pointing upwards vs. downwards). The training phase stimuli were each composed of one Rule attribute which was used to classify the stimuli, one attribute that was perfectly correlated (PC) with the Rule attribute, and four family resemblance (FR) attributes, each of which had a .75 correlation with the Rule attribute. Stimuli were divided into two categories, as described in Table 1. The Rule attribute was perfectly predictive of category membership. Participants were explicitly told about and shown the Rule attribute prior to the training phase. However, they were not given any information about the PC or the FR attributes.

During the test phase, participants were presented with an equal number of old and new items. The 48 new items were created by replacing the value of one or many attributes by values diagnostic of the opposite category. These modifications yielded six different new item types. Modifying the value of one or two FR attributes created the 1FR and 2FR new item types. Modifying the value of the

Table 1: Categorical structure for training phase (Experiments 1 and 2)

Category membership	*Attribute*					
	Rule	*PC*	*FR1*	*FR2*	*FR3*	*FR4*
Category A	0	0	0	0	0	1
	0	0	0	0	1	0
	0	0	0	1	0	0
	0	0	1	0	0	0
Category L	1	1	1	1	1	0
	1	1	1	1	0	1
	1	1	1	0	1	1
	1	1	0	1	1	1

	Attributes					
	Rule	*PC*	*FR1*	*FR2*	*FR3*	*FR4*
Original OLD item	0	0	0	0	0	1
1FR test item	0	0	1	0	0	1
2 FR test item	0	0	1	1	0	1
PC test item	0	1	0	0	0	1
PC+1FR test item	0	1	1	0	0	1
PC+2FR test item	0	1	0	1	1	1

Figure 2: Sample test items created using a single training logical exemplar (Experiment 1). Shaded cells represent the attributes that were reversed.

PC attribute created the PC new item type. Finally, modifying the value of the PC attribute as well as that of one or two FR attributes yielded the PC+1FR and PC+2FR new item types. Figure 2 shows the 5 new item types created from one old exemplar.

Procedure Participants were randomly assigned to one of six experimental conditions. In order to create the experimental conditions, we assigned different physical properties to each of the six attributes of the training phase stimuli. Because it was not feasible to test all 720 possible permutations, we selected six with the constraint that each shape appeared once as the Rule, as the PC attribute, and as each of the four FR attributes across all six conditions. Table 2 shows the assignment of physical properties to attributes in each of the six conditions.

The experiment lasted approximately 60 minutes. All training trials proceeded as follows. First, a fixation point appeared in the center of the screen for 500 ms. A hexagonal display stimulus was then presented. Participants searched for the Rule attribute and used this perfectly predictive attribute to classify the display stimulus by selecting the appropriate letter on the keyboard (A or L). The position of each of the six attributes was randomly determined on each trial.

Feedback pertaining to accuracy was provided. When the participants were correct, the phrase ‘‘correct response’’ appeared for 750 ms; when they were incorrect, the participants heard a short tone and the phrase ‘‘incorrect response’’ appeared for 750 ms. The interstimulus interval was 1000 ms. Participants were instructed to maintain a cumulative minimal accuracy of 95% correct responses and a maximal cumulative average RT of 1000 ms throughout the experiment. The actual cumulative averages were communicated after every fourth training block and participants received warnings if the criteria were not respected. Each training item was presented once within a block. Participants achieved 80 training blocks, for a total of 640 trials.

Following completion of the training phase, participants were instructed to continue categorizing the stimuli into the “A” and “L” categories as quickly and accurately as possible. Participants were not told about the inclusion of new items. Test trials were identical to training trials except

Table 2: Experimental conditions (Experiments 1 and 2)

Rule condition	*Attribute*					
	Rule	*PC*	*FR1*	*FR2*	*FR3*	*FR4*
1 (Texture)	A	B	C	D	E	F
2 (Color)	B	F	D	C	A	E
3 (Orientation)	C	D	E	F	B	A
4 (Size)	D	A	F	E	C	B
5 (Filling)	E	C	A	B	F	D
6 (Direction)	F	E	B	A	D	C

Note: A = heart; B = diamond; C = square; D = spade; E = circle; F = arrow.

for the absence of feedback. Participants achieved 96 test trials. A different logical exemplar (either old or new) was chosen randomly on each trial.

Results

Training Participants' overall accuracy was 97.5% (SD = 3.5%) and their response time average for correct trials was 560ms (SD = 138ms). To analyze the effects of training on accuracy and response times (RTs) for each condition, the data collected from the 80 training blocks were separated into 8 groups of ten blocks each. The participants' percentage correct and response times were subjected to separate 6 (Condition: 1 through 6) x (8) (Block: 8 levels) analyses of variance (ANOVAs). The accuracy data showed no significant main effect or interaction (all ps >.05), but produced a tendency towards a main effect for Block (p = .051), suggesting a slight increase in accuracy from Block 1 (96.4%) to Block 8 (97.8%). As for RTs, both main effects were significant (Block: p = .03; Condition: p = .01). *Post-hoc* comparisons showed a significant decrease from Block 1 (657ms) to Block 8 (541ms), as well as generally slower response times for the “orientation” rule condition (673 ms). The interaction did not reach significance (p > .05).

Test Participants' overall accuracy was 97.1% (SD = 6.1%) and their RT average for correct trials was 544ms (SD = 171ms). The participants' accuracy scores and RTs were subjected to separate 6 (Condition: 1 through 6) x (6) (Item type: OLD, 1FR, 2FR, PC, PC+1FR, PC+2FR) ANOVAs.

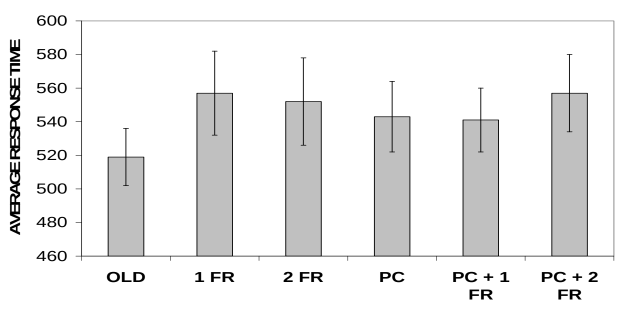

Figure 3: Average response times for the six item types used in the Experiment 1 test phase.

For accuracy data, no main effect or interaction was significant (all *p*s > .05). As for RTs (Figure 3), the only significant main effect was that of Item type (p = .01). *Post-hoc* comparisons showed that RTs for old items (519 ms) were significantly lower than those for all new item types (1FR: 557ms; 2FR: 552 ms; PC: 543 ms; PC+1FR: 541 ms; PC+2FR: 557ms; all *p*s < .05), but that RTs did not significantly differ among new item types (all *p*s > .05). No interaction was significant. Surprisingly, the results show that violating the correlations among within-category attributes slowed down performance for all items types. Thus, participants seemed to have incidentally learned the relationship between the Rule attribute and the PC attribute, as well as the relationship between the Rule attribute and the different FR attributes.

Experiment 2

The question behind the second experiment was as follows: given that the visual search categorization task allows implicit, perceptual learning of within-category correlations, can this knowledge be later used explicitly to classify incomplete stimuli? Recall that participants in Giguère et al.'s (2007) study had little success in categorizing stimuli lacking the Rule attribute except when a highly salient and perfectly correlated attribute (i.e., color) was present. We therefore expected that participants assigned to a rule condition in which the PC attribute was color would exhibit above chance accuracy rates when the Rule attribute was removed in the transfer phase. Because participants were required to visually seek out the Rule attribute, we expected that they might be more likely to incidentally learn non-rule FR and PC attributes which were less salient than color.

Method

Participants 72 undergraduate students from Carleton University participated in the experiment. They received course credit as compensation for their time.

Material The stimuli and categorical structure used for the training phase were identical to those of Experiment 1. At test, participants were only presented with 152 new items, which were incomplete versions of training items created by deleting one or many of their attributes. The Rule attribute was absent on all test trials. Deleting the value of the Rule the PC+1FR and PC+2FR test item types. Figure 4 shows the different test items created from one old exemplar.

	Attributes					
	Rule	*PC*	*FR1*	*FR2*	*FR3*	*FR4*
Original training item	0	0	0	0	0	1
RULE test item		0	0	0	0	1
1FR test item		0		0	0	1
2 FR test item		0		0		1
PC test item			0	0	0	1
PC+1FR test item			0	0		1
PC+2FR test item					0	1

Figure 4 Sample test items created using a single training logical exemplar (Experiment 2). Shaded cells represent the attributes that were reversed.

Procedure The training procedure was identical to that of Experiment 1. Following completion of the training phase, participants were instructed to continue categorizing the displays into the "A" and "L" categories using the information that remained on the display. They were also warned that the Rule, as well as various other attributes, would no longer appear. They were encouraged to perform as accurately as possible despite the removal of the Rule attribute. It was also noted that speed was no longer important. Test trials were identical to training trials, apart from the absence of feedback. Participants achieved 152 test trials.

Results

Training Participants' overall accuracy was 97% (*SD* = 2.5%) and their RT average for correct trials was 546 ms (*SD* = 87 ms). Both measures were submitted to separate 6 (Condition: 1 through 6) x (8) (Block: 8 levels) ANOVAs. The accuracy data produced a significant main effect for Block (p < .001). The percentage correct increased from 96% in the first block to 98% in the last. For RTs, there was a Block x Condition interaction (p = .033). It was found that participants' RTs became faster with practice (the average Block 1 RT was 624 ms vs. 520 ms for the last block), but that certain conditions still elicited faster responses times than others at the end of training. The rule condition that yielded the fastest average Block 8 RTs was that of color (472 ms) and, once again, the condition that yielded the slowest was that of orientation (647 ms).

Test Because speed was not stressed during the test phase, the RT data are not reported here. Participants' overall mean accuracy was 53.0% (*SD* = 8.4%), reflecting an overall near chance performance. Participants' percent accuracy were subjected to a 6 (Condition: 1 through 6) x (6) (Item type: Rule, 1FR, 2FR, PC, PC+1FR, PC+2FR) ANOVA. The results are shown in Figure 5.

A significant Condition x Item type interaction was found (p = .03). An inspection of Figure 5 shows that performance appears to be above chance for certain item types for which the PC attribute was not removed (Rule/1FR/2FR types). This would indicate that some of participants were able to incidentally learn perfect within-category correlations and use them to classify the stimuli in the absence of the Rule attribute. To verify this impression, one-sample t-tests were conducted within each condition to determine if the performance for each item type significantly differed from random performance. The significance level was set at α = .05 and the p-values were adjusted using the Bonferroni Step-down correction. When the PC attribute was color, three item types produced responding beyond chance: Rule,

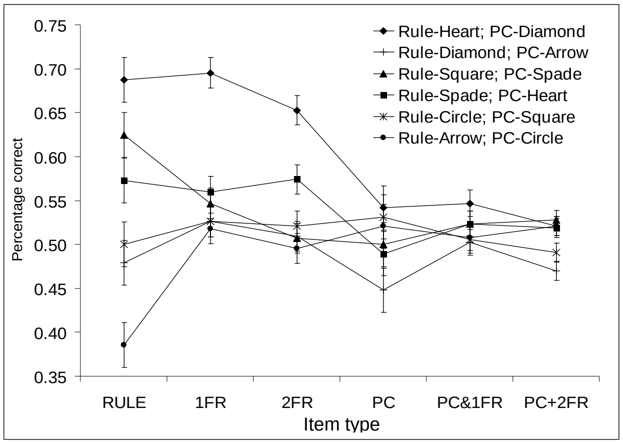

Figure 5: Participants' percentage of correct classifications for item type by condition.

1FR, and 2FR. When the PC attribute was size, one item type produced responding beyond chance, namely Rule. Finally, when the PC attribute was texture, two item types produced responding beyond chance: 1FR and 2FR. No other test reached significance. These results replicate and extend Giguère et al.' (2007) results. Some participants incidentally learned within-category correlations and used this knowledge to categorize stimuli from which the Rule attribute had been removed.

Discussion

The results of Experiment 1 show that all the non-rule attributes of category members, whether PC or FR, were learned incidentally. These results cannot be explained by a speed-accuracy trade-off, because the accuracy level did not differ across old vs. new item types. They cannot be explained either by a generalized caution effect, because caution would also have affected latencies for old items as well as the new items. Finally, these results cannot be explained by an exemplar learning strategy, yielding an advantage for training items, because the same exact physical configuration of shapes was never seen twice during training, making it impossible for participants to memorize physical exemplars.

Results from Experiment 2, while not as unequivocal, replicated Giguère et al.'s (2007) earlier results. They show that, except in rare conditions, participants cannot access the implicit knowledge that they have of the within-category attribute correlations to consciously categorize the stimuli in the absence of the rule attribute. Taken together, the results of the two experiments suggest that classification tasks such as that used in Experiment 2, in Giguère et al.'s study and in Chin-Parker and Ross' (2002; 2004) studies may have been underestimating the amount of implicit, perceptual knowledge that is acquired about within-category correlations.

Extensions of this work will focus on the content of postexperimental interviews to determine the relationship between performance in the two tasks used here and verbalizable knowledge about within-category correlations. One would expect little or no relationship in a task such as that of Experiment 1 and a strong relationship in a task such as that of Experiment 2.

Acknowledgments

This research was made possible by grants awarded to the second and fourth author by the Natural Sciences and Engineering Research Council of Canada (NSERC).

References

Anderson, J. R., & Fincham, J. M. (1996). Categorization and sensitivity to correlation. *Journal of Experimental Psychology: Learning, Memory, and Cognition. Vol. 22,* 259-277.

Anderson, A. L., Ross, B. H., & Chin-Parker, S. (2002). A further investigation of category learning by inference. *Memory and Cognition, 30,* 119-128.

Bruner, J., Goodnow, J.J., & Austin, G.A. (1956). *A study of thinking.* New York: Wiley.

Chin-Parker, S., & Ross, B. H. (2002). The effect of category learning on sensitivity to within-category correlations. *Memory and Cognition, 30,* 353-362.

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.*

Giguère, Lacroix, & Larochelle (2007). Learning the correlational structure of stimuli in a one-attribute classification task. *European Journal of Cognitive Psychology, 19,* 457-469

Lacroix, G. L., Giguère, G., & Larochelle, S. (2005). The origin of exemplar effects in rule-driven categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31,* 272-288.

Lassaline, M. E., & Murphy, G. L. (1996). Induction and category coherence. *Psychonomic Bulletin and Review, 3,* 95-99.

Nosofsky, R. M., Palmeri, T. J., & McKinley, S. C. (1994). Rule-plus-exception model of classification learning. *Psychological Review*, *101*, 53-79.

Regehr, G., & Brooks, L. R. (1993). Perceptual manifestations of an analytic structure: The priority of holistic individuation. *Journal of Experimental Psychology: General, 122,* 92-114.

Smith, J. D., Redford, J. S., Gent, L. C., & Washburn, D. A. (2005). Visual search and the collapse of categorization. *Journal of Experimental Psychology: General, 134,* 443-460.

Thomas, R. (1998). Learning correlations in categorization tasks using large, ill-defined categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition. Vol. 24,* 119-143.