

This article was downloaded by:[Carleton University]  
[Carleton University]

On: 16 April 2007

Access Details: [subscription number 769425782]

Publisher: Psychology Press

Informa Ltd Registered in England and Wales Registered Number: 1072954

Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK



## European Journal of Cognitive Psychology

Publication details, including instructions for authors and subscription information:  
<http://www.informaworld.com/smpp/title-content=t713734596>

### Learning the correlational structure of stimuli in a one-attribute classification task

First Published on: 29 August 2006

To cite this Article: 'Learning the correlational structure of stimuli in a one-attribute classification task', *European Journal of Cognitive Psychology*, 19:3, 457 - 469

To link to this article: DOI: 10.1080/09541440600926716

URL: <http://dx.doi.org/10.1080/09541440600926716>

PLEASE SCROLL DOWN FOR ARTICLE

Full terms and conditions of use: <http://www.informaworld.com/terms-and-conditions-of-access.pdf>

This article maybe used for research, teaching and private study purposes. Any substantial or systematic reproduction, re-distribution, re-selling, loan or sub-licensing, systematic supply or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The accuracy of any instructions, formulae and drug doses should be independently verified with primary sources. The publisher shall not be liable for any loss, actions, claims, proceedings, demand or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

© Taylor and Francis 2007

## Learning the correlational structure of stimuli in a one-attribute classification task

Gyslain Giguère

*Laboratoire d'Études en Intelligence Naturelle et Artificielle, Université du Québec à Montréal, Montreal, Canada*

Guy L. Lacroix

*Department of Psychology, Carleton University, Ottawa, Canada*

Serge Larochelle

*Département de psychologie, Université de Montréal, Montreal, Canada*

In category learning experiments, participants typically do not learn within-category correlations unless the composition of the categories or the task demands compel them to do so. To determine if correlations among attributes could be learned without explicitly focusing the participants' attention on them, a task was designed that allowed stimuli to be classified on the basis of a single perfectly predictive attribute. Each training stimulus also included attributes that were either perfectly or partly correlated with the rule attribute. Then, in a test phase, the impact of eliminating the rule attribute on classification was evaluated. The experiment showed that some of the attributes that were perfectly correlated with the rule attribute were learned. These attributes could be used to classify the test exemplars even though the rule attribute had been removed. This experiment provides evidence that within-category correlations can be learned incidentally during classification tasks.

Through the years, results from category construction tasks,<sup>1</sup> which involve asking participants to sort exemplars into groups or categories that seem natural, have repeatedly shown that the participants' preferred strategy is to choose one salient dimension and divide the exemplars accordingly, although the attributes that enter in the composition of the exemplars

<sup>1</sup>Also known as "sorting" tasks.

Correspondence should be addressed to Serge Larochelle, Département de psychologie, Université de Montréal, C.P. 6128 Succ. Centre-Ville, Montréal, QC, Canada H3C 3J7. E-mail: serge.larochelle@umontreal.ca

This research was supported by grants awarded to the third author by the Université de Montréal, as well as from the Natural Sciences and Engineering Research Council.

afford more complex, multiattribute (family resemblance) sorting (Ahn & Medin, 1992; Medin, Wattenmaker, & Hampson, 1987; Regehr & Brooks, 1995). Here, the prevailing inclination for one-attribute sorting can seem rather unsurprising, because there is no response-driven performance feedback to force the participants to learn multiple attribute-category label associations.

Results from category learning experiments have also consistently led researchers to the conclusion that participants typically adopt a hypothesis-testing strategy which involves successively trying simple, one-attribute rules.<sup>2</sup> This phenomenon can be observed clearly in category induction tasks, in which the presence of feedback compels participants to try this strategy in order to “solve” the categorisation problem and perform accurately (Nosofsky, Palmeri, & McKinley, 1994; see also Erickson & Kruschke, 1998; Love, Medin, & Gureckis, 2004). Nevertheless, when the categories to be learned do not allow the one-dimension strategy, participants must pay attention to many attributes in order to successfully categorise stimuli, thus longitudinally creating exemplar memories. The resulting representation will generally be incomplete, however, as the memories only encompass the attributes involved in making categorical judgements (i.e., the diagnostic attributes), instead of complete stimuli (Chin-Parker & Ross, 2004).

The preceding conclusion is supported by another 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. Regehr and Brooks (1993) 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 participants had learned the relationships between the rule and the nonrule attributes. Lacroix et al. (2005) replicated this result and showed that the two attributes not specified in the rule could be modified without consequence for the categorisation results. They concluded 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 led to the conclusion that, generally, it is the attention that is paid to the stimulus attributes that determines what is learned during the categorisation process, and that participants are not naturally inclined to

---

<sup>2</sup> For a review, see Murphy (2002).

learn relationships between attributes unless their attention is focused on them.

In fact, the only task that has successfully shown consistent learning of within-category attribute correlations is inference learning (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 the category label. This type of learning is one of the rare ways of making participants override their propensity for one-attribute classifications, observed in category construction tasks, in favour of multiattribute classifications (see also Lassaline & Murphy, 1996). Asking participants to generate inferences concerning the relationships between attributes therefore appears to be a powerful way of learning within-category attribute associations, because the participants' attention is explicitly focused on these relations. This allows learning to take place and subsequently a more successful recall of the attribute associations that can influence categorical decisions.

To summarise, it can be argued that the category learning literature has produced two consistent, robust findings. First, most participants' default strategy in category learning experiments is to seek simple and very often one-attribute rules to classify the stimuli. Second, participants do not deviate from this strategy, nor do they typically learn the correlations among the attributes comprising the stimuli, unless they are compelled to do so by the composition of the categories or the task demands.

## A ONE-ATTRIBUTE CLASSIFICATION TASK

The goal of this study was to explore if there existed conditions that could lead to the learning of within-category correlations in a classification task without explicitly focusing the participants' attention on these relations. We designed a classification task, which allowed creature-like stimuli to be categorised on the basis of one attribute that was perfectly predictive of category membership. Participants were told to use this "rule" attribute to classify the stimuli. Our paradigm thus utilised the participants' propensity to categorise exemplars using simple one-attribute rules, while specifying the attribute to be used. Training stimuli also included attributes that were either perfectly or partly correlated with the rule attribute. In the test phase of the experiment, participants were shown test items that were highly similar to training items. However, the rule attribute was removed and the participants were instructed to classify the creatures in the same two categories as before. This procedure created a situation in which the participants were made to show any additional knowledge gained about the creatures in the training phase, because only the correlated attributes not mentioned in the training

instructions remained. In a subsequent test phase, the perfectly correlated attribute was also removed, leaving only the partly correlated attributes as a basis for categorisation.

There are several reasons why Lacroix et al. (2005) failed to find evidence of learning the nonrule attributes in a rule-based categorisation task. The rules used in their study were quite complex, involving many attributes, thereby placing a burden on attention that may have prevented the consideration of nonrule attributes. Moreover, the nonrule attributes were located at the periphery of the creatures (tail and neck) and were totally uninformative about category membership. Finally, the abstract attributes were implemented using idiosyncratic physical features so that it is not clear whether the participants considered these features as the embodiment of a unique abstract attribute value (e.g., straight tail),<sup>3</sup> especially in the case of nonrule attributes. The situation is quite different in the present experiment: The rule is very simple (it involves a single attribute), each value of a given attribute has the same physical embodiment (e.g., all red creatures have the same red colour), and all nonrule attributes are correlated with the rule attribute (and are therefore informative of category membership). Finally, the position of the rule and nonrule attributes varies within the creatures: Some attributes span the entire stimulus (e.g., colour), while others are “local” (e.g., number of legs).

The rationale of the experiment is as follows. Neither the composition of the stimuli, nor the task demands required participants to pay attention to any attribute except the rule attribute. Above-chance performance in the test phases would therefore yield evidence that participants can learn within-category correlations in a categorisation task that does not require attention to be paid to multiple attributes. We hypothesised that participants would be more likely to learn associations between the rule attribute and the most strongly correlated attributes, especially if the latter spanned the entire stimulus.

## METHOD

### Participants and materials

Ninety-six students at the Université de Montréal participated in the study. Participants were randomly assigned to one of eight conditions. They received \$3 as compensation for their time if they had 40 training trials and \$5 if they had 160.

---

<sup>3</sup> Yamauchi and Markman (2000; see also Markman & Maddox, 2002) found that increasing the number of different manifestations of abstract features made classification tasks more difficult.

The 10 training phase stimuli were drawings of fictional creatures composed of seven binary attributes: colour (green vs. yellow), texture (composed of lines vs. pluses), body type (rounded vs. angular), dorsal pattern (spots vs. stripes), head shape (oval vs. D-shaped), tail type (straight vs. cane-shaped), and number of legs (two vs. four). Two examples of the training stimuli are given in the left panel of Figure 1. The top part of Table 1 gives an abstract description of the 10 training stimuli that were divided in two categories. Each animal was composed of one “rule attribute”, which determined category membership, one attribute that was perfectly correlated (PC) with the rule attribute, and five family resemblance (FR) attributes, each of which had a .80 correlation with the rule attribute. In the training phase, the rule attribute was given to the participants as the basis for classifying the creatures. The PC and FR attributes were never mentioned.

Although all the stimuli were built using the abstract structure in Table 1, two different sets of exemplars were created to vary the type of rule and PC attributes used during training. In the “global” condition, both the rule and the PC attributes were realised using dimensions that spanned the entire stimulus. They were colour and texture. The remaining attributes—body type, head shape, dorsal pattern, number of legs, and tail type—served as FR attributes. In the “local” condition, the rule attribute and the PC attribute were realised using dimensions that were confined to a limited part of the stimulus. They were body type and number of legs. The FR attributes were colour, texture, head shape, dorsal pattern, and tail type.



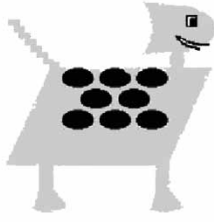
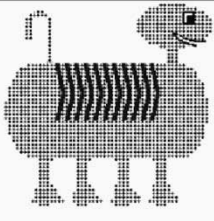
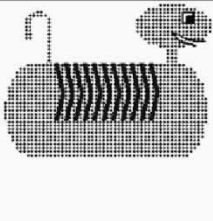

	Training stimuli	Corresponding PC&FR stimuli	Corresponding FR stimuli
Global/ High saliency condition		 <i>Texture removed</i>	 <i>Texture and colour removed</i>
Local/ Low saliency condition		 <i>Legs removed</i>	 <i>Legs and body removed</i>

Figure 1. Black and white examples of the stimuli used in the experiment.

TABLE 1  
Abstract description of the stimuli used in the experiment

<i>Category membership</i>													
		<i>A</i>					<i>B</i>						
<i>RULE</i>	<i>PC</i>	<i>FR1</i>	<i>FR2</i>	<i>FR3</i>	<i>FR4</i>	<i>FR5</i>	<i>RULE</i>	<i>PC</i>	<i>FR1</i>	<i>FR2</i>	<i>FR3</i>	<i>FR4</i>	<i>FR5</i>
Training items													
0	0	0	0	0	0	1	1	1	1	1	1	1	0
0	0	0	0	0	1	0	1	1	1	1	1	0	1
0	0	0	0	1	0	0	1	1	1	1	0	1	1
0	0	0	1	0	0	0	1	1	1	0	1	1	1
0	0	1	0	0	0	0	1	1	0	1	1	1	1
Test items													
~	0 or ~	0	0	0	0	1	~	1 or ~	1	1	1	1	0
~	0 or ~	0	0	0	1	0	~	1 or ~	1	1	1	0	1
~	0 or ~	0	0	1	0	0	~	1 or ~	1	1	0	1	1
~	0 or ~	0	1	0	0	0	~	1 or ~	1	0	1	1	1
~	0 or ~	1	0	0	0	0	~	1 or ~	0	1	1	1	1
~	0 or ~	0	0	0	1	1	~	1 or ~	1	1	1	0	0
~	0 or ~	0	0	1	1	0	~	1 or ~	1	1	0	0	1
~	0 or ~	0	1	1	0	0	~	1 or ~	1	0	0	1	1
~	0 or ~	1	1	0	0	0	~	1 or ~	0	0	1	1	1
~	0 or ~	1	0	0	0	1	~	1 or ~	0	1	1	1	0
~	0 or ~	0	0	1	0	1	~	1 or ~	1	1	0	1	0
~	0 or ~	0	1	0	1	0	~	1 or ~	1	0	1	0	1
~	0 or ~	1	0	1	0	0	~	1 or ~	0	1	0	1	1
~	0 or ~	1	0	0	1	0	~	1 or ~	0	1	1	0	1
~	0 or ~	0	1	0	0	1	~	1 or ~	1	0	1	1	0

The global and local conditions were separated depending on the saliency of the rule attribute: high (colour and body type) vs. low (texture and number of legs). Attribute saliency was determined on the basis of results obtained in a previous category construction experiment involving similar stimuli (Lacroix & Larochelle, 2000).

Two types of test item were created. The stimuli of the first type were composed of all the attributes of the training stimuli except for the rule attribute. Since the PC and FR attributes remained unchanged, they will be called the PC&FR stimuli. Four different PC&FR stimulus sets were created, depending on whether the rule attribute deleted was colour (global, high saliency), texture (global, low saliency), body type (local, high saliency), or number of legs (local, low saliency). The second type of test item will be called the FR stimuli because both the rule and the PC attributes were removed, leaving only the five FR attributes. The FR condition required two more stimulus sets: one for the participants with a global rule attribute and one for the participants with a local rule attribute. There were 30 test stimuli of each type: 10 test items had configurations of FR attributes that had been seen in the training phase and 20 items had a new configuration. Sample stimuli for the PC&FR and FR test phases are shown in the middle and right panels of Figure 1, respectively.

## Procedure

The participants were tested individually, using the program MEL Professional v.2.01 (Schneider, 1989). They were randomly assigned to one of four rule conditions and to one of two training length conditions: 4 blocks or 16 blocks. The experiment was conducted in three phases. In the training phase, the participants were instructed to classify the creatures in two different families using the one-attribute rule provided. They were also asked to answer as quickly as possible while avoiding mistakes and to stay concentrated even if they found the task easy. Each block was composed of the 10 training items presented in random order. All training trials proceeded as follows. First, a fixation point appeared in the centre of the screen for 1500 ms. An exemplar was then presented and participants had to classify the stimulus by selecting the appropriate letter on the keyboard. All items appeared on a black background. Feedback pertaining to accuracy was provided. When the participants were correct, the phrase "correct response" appeared for 2000 ms; when they were incorrect, the participants heard a short tone and the phrase "incorrect response" appeared for 2000 ms. The interstimulus interval was 1000 ms.

The PC&FR stimuli were used in the first part of the test phase. Participants were told that they would see members of the same two families

of creatures, but that the creatures would appear without the rule attribute, which was appropriately named. Their task was to try to classify the creature in the same families as before. They were told to do their best, to respond spontaneously, and to guess if necessary. The 30 test items that were without the rule attribute were presented once each. Therefore, these items could be classified on the basis of any combination of PC and/or FR attributes. The procedure was identical to that used in the training phase, except for the absence of feedback. The second part of the test phase was identical to the preceding one, except that it involved the FR stimuli, both the rule and the PC attributes having been removed from the items. Hence, correct classification required knowledge of the FR attributes. Finally, postexperimental interviews were conducted in which participants were asked how they had proceeded to classify the stimuli in both test phases.

## RESULTS

### Training phase

To evaluate potential differences between conditions due to the two sets of training stimuli that were used, an analysis of variance was performed on the percentage of correct classifications obtained in this phase of the experiment. This analysis involved a  $2 \times 2$  design with two between-subjects factors: feature type used as the rule (global vs. local) and training length (40 trials vs. 160 trials). The average percentage of correct classification was 98.1% ( $SD = 0.02\%$ ). The interaction between the factors was not significant,  $F(1, 92) = 0.25$ ,  $MSE = 0.00$ ,  $p = .62$ , nor was the main effect of feature type used as the rule,  $F(1, 92) = 0.37$ ,  $MSE = 0.0001$ ,  $p = .55$ , or of training length,  $F(1, 92) = 2.73$ ,  $MSE = 0.002$ ,  $p = .1$ . Thus, participants were highly accurate in classifying all stimuli.

### Test phases

An analysis of variance was performed on the percentage of correct classifications. It involved a  $2 \times 2 \times 2 \times (2)$  design with three between-subjects factors: training length (40 trials vs. 160 trials), feature type used as the rule (global vs. local), saliency of the rule attribute (high vs. low); and one within-subject factor: test phase (PC&FR vs. FR).

Classification accuracy for all conditions is presented in Figure 2. The top portion shows the data obtained after 40 trials of practice and the bottom portion shows the data obtained after 160 trials. As can be observed, general performance was poor (55.1%). The interaction involving all four factors was not significant,  $F(1, 80) = 1.8$ ,  $MSE = 0.04$ ,  $p = .18$ . However, there was

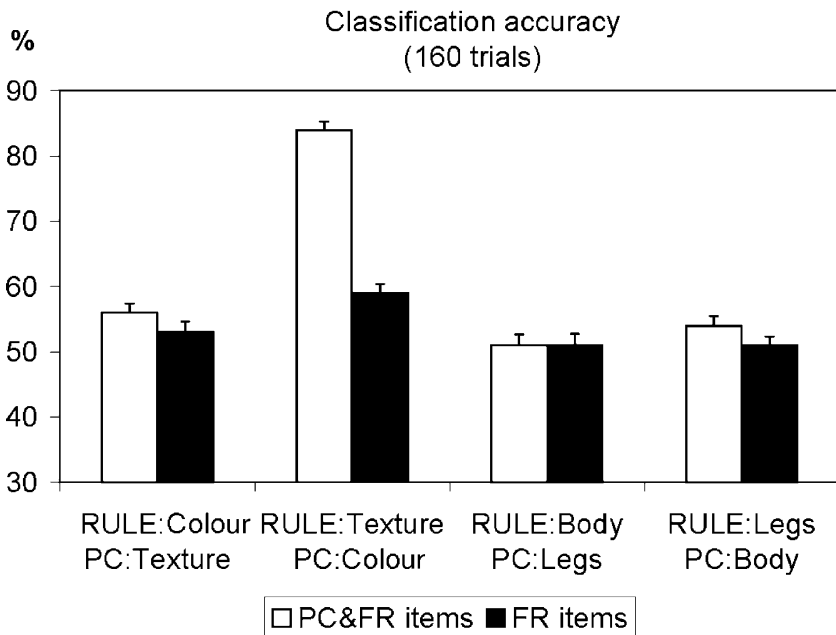



Figure 2. Average test classification accuracy in the experiment.

a three-way interaction involving feature type, saliency, and test phase,  $F(1, 80) = 10.9$ ,  $MSE = 0.24$ ,  $p = .01$ . Decomposition of this interaction showed that the interaction between feature type and saliency was significant for PC&FR items,  $F(1, 80) = 16.6$ ,  $MSE = 0.68$ ,  $p < .001$ , but not for FR items,  $F(1, 80) = 0.96$ ,  $MSE = 0.02$ ,  $p = .33$ . Decomposition of the interaction between feature type and saliency for the PC&FR items showed that when the feature type was global, there was a significant difference between the high and low saliency rule attribute conditions,  $F(1, 80) = 21.1$ ,  $MSE = 0.87$ ,  $p < .001$ , but no such difference was found when the feature type was local,  $F(1, 80) = 1.2$ ,  $MSE = 0.05$ ,  $p = .28$ . Simply put, when the rule attribute was texture and the PC attribute was colour, participants' classification accuracy was much greater (81%) than in all other conditions (52%) involving PC&FR items. With FR items, performance was uniformly poor (52%). The training length factor did not interact with any other factor.

## DISCUSSION

The present study required participants to classify stimuli using a one-attribute rule, but offered them the opportunity to show deeper, explicit knowledge of the stimuli and category structures via the use of test phases. More precisely, the goal of this experiment was to determine whether knowledge of within-category correlations could be learned and become explicit in a simple one-attribute rule classification task.

According to Markman and Ross (2003), "information that is necessary to complete a task is much more likely to be learned than any other information that is present during learning, but is not necessary for the successful completion of the task. Furthermore, because people are miserly in their allocation of cognitive effort, they may not learn much additional information" (p. 595). By design, our one-attribute rule classification paradigm did not require participants to focus on attributes other than the rule attribute in order to successfully accomplish the task. Thus, participants had absolutely no incentive to discover the relations between other attributes and category membership as can be the case for induction and inference learning paradigms. Yet, our experiment leads to the conclusion that some knowledge of the within-category correlations can be gained, by contrast with what is typically found in category construction experiments, and despite the known propensity for conservatism in cognitive resource allocation (Payne, Bettman, & Johnson, 1993).

When texture was the rule attribute and colour was a perfectly correlated attribute, 81% of the participants were able to classify the transfer items although the rule attribute had been removed. No other training combination involving rule and PC attributes led to performance above that

predicted by chance in the first transfer phase. In the second transfer phase, when both the rule and PC attributes had been removed, chance performance was again found. This result held even when colour was a FR attribute. Postexperimental interviews revealed that that 22 of the 48 participants chose this attribute to classify the stimuli in the PC&FR and FR test phases combined. Yet, the accuracy rates still reflected chance performance. Therefore, participants were successful in correctly associating a given colour and the appropriate category *only* when this attribute was perfectly diagnostic of category membership.

Hence, our results show that participants did learn within-category correlations when certain criteria were met. First, in order to be noticed, the within-category correlations needed to be perfect. Participants were unable to classify the stimuli on the sole basis of the FR attributes, even though each one had a fairly high correlation (.80) with the rule. Another result is that participants could only categorise the test stimuli correctly when colour was the PC attribute. The fact that performance was at chance level in all other cases confirms that participants did not explicitly know the association between the other attribute values and the corresponding categories. Thus, high attribute saliency was a second necessary condition for the effect to obtain.

The most parsimonious explanation for the effect is that the lower cognitive demands of the task (i.e., using a single attribute to sort a small set of exemplars deterministically) allowed the participants to give sporadic attention to the nonrule attributes. In the context of selective attention research, it has been shown that participants' attention will often "slip" to nontargeted information or cues when adequate methodological steps are not taken to prevent this from happening (Lachter, Forster, & Ruthruff, 2004). In our experiment, participants were asked to use the rule attribute to categorise the stimuli, but they were free to notice the other attributes. Indeed, it was the goal of our experiment to explore if such incidental learning would take place. Thus, it is plausible that these occasional attentional slippages were sufficient to allow the participants to discover the relation between the rule attribute and another attribute if the latter was perfectly correlated with the former.

In further research, it would be interesting to determine more precisely what parameters allow such incidental learning of nonrule attributes to take place. For instance, an attribute cue validity higher than 80% might lead to successful learning without necessarily reaching 100%. It would also be intriguing to determine if the use of multiattribute rules, such as those used in Lacroix et al. (2005; Regehr & Brooks, 1993), would diminish the probability that perfectly diagnostic nonrule attributes be learned, due to the greater cognitive demand of applying a complex rule. Also, we know that colour and texture were not viewed as integral dimensions or as a single

feature when they were perfectly correlated. Had this been the case, one would have obtained similar results when texture was the PC attribute and colour was the rule attribute. This result suggests that a correlated attribute has to possess a very high level of saliency in order to be incidentally noticed and learned. Although evidence of nonrule learning was restricted to colour in the present experiment, one can imagine other global attributes, such as body size for example, to be salient enough to also produce nonrule attribute learning in a single-attribute classification task.

Original manuscript received October 2005

Revised manuscript received July 2006

First published online 29 August 2006

## REFERENCES

- Ahn, W. K., & Medin, D. L. (1992). A two-stage model of category construction. *Cognitive Science*, *16*, 81–121.
- Anderson, A. L., Ross, B. H., & Chin-Parker, S. (2002). A further investigation of category learning by inference. *Memory and Cognition*, *30*, 119–128.
- 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.
- Erickson, M. A., & Kruschke, J. K. (1998). Rules and exemplars in category learning. *Journal of Experimental Psychology: General*, *127*, 107–140.
- Lachter, J., Forster, K. I., & Ruthruff, E. (2004). Forty-five years after Broadbent: Still no identification without attention. *Psychological Review*, *111*, 880–913.
- 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.
- Lacroix, G. L., & Larochelle, S. (2000). *Learning family resemblance attributes in a one-dimension rule classification task*. Unpublished manuscript.
- Lassaline, M. E., & Murphy, G. L. (1996). Induction and category coherence. *Psychonomic Bulletin and Review*, *3*, 95–99.
- Love, B. C., Medin, D. L., & Gureckis, T. M. (2004). SUSTAIN: A network model of category learning. *Psychological Review*, *111*, 309–332.
- Markman, A. B., & Maddox, W. T. (2002). Classification of exemplars with single and multiple feature manifestations: The case of relevant dimension variation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *29*, 107–117.
- Markman, A. B., & Ross, B. H. (2003). Category use and category learning. *Psychological Bulletin*, *129*, 592–613.
- Medin, D. L., Wattenmaker, W. D., & Hampson, S. E. (1987). Family resemblance, concept cohesiveness, and category construction. *Cognitive Psychology*, *19*, 242–279.
- Murphy, G. L. (2002). *The big book of concepts*. Cambridge, MA: MIT Press.
- Nosofsky, R. M., Palmeri, T. J., & McKinley, S. C. (1994). Rule-plus-exception model of classification learning. *Psychological Review*, *101*, 53–79.

- Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The adaptive decision maker*. New York: Cambridge University Press.
- 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.
- 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*, 347–363.
- Schneider, W. (1989). Micro experimental laboratory: An integrated system for IBM PC compatibles. *Behavior Research Methods, Instruments and Computers*, *20*, 206–217.
- Yamauchi, T., & Markman, A. B. (2000). Learning categories composed of varying instances: The effect of classification, inference and structural alignment. *Memory and Cognition*, *24*, 64–78.