UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title
Categorization vs. Inference: Shift in Attention or in Representation?

Permalink
https://escholarship.org/uc/item/15p5w118

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 27(27)

ISSN
1069-7977

Author
Niyogi, Sourabh

Publication Date
2005

Peer reviewed

Categorization vs. Inference: Shift in Attention or in Representation?

Håkan Nilsson (hakan.nilsson@psyk.uu.se)
Department of Psychology, Uppsala University
SE-741 42, Uppsala, Sweden

Henrik Olsson (henrik.olsson@psyk.uu.se)
Department of Psychology, Uppsala University
SE-741 42, Uppsala, Sweden

Abstract

Recently it has been found that people that learn through inference create qualitatively different cognitive representations than those who learn through categorization. The present study addresses the question of whether the findings supporting this claim generalize to a design where both learning tasks have a probabilistic relation between each stimulus cue and the category label. It was shown that participants in the categorization condition learned faster than participants in the inference condition. Further, participants in the inference condition did not rely on prototypical values when making one-cue categorizations. The results suggest that shifts in attention must be considered as a viable explanation of some of the results in studies that investigate differences between inference and categorization.

Introduction

It has been argued that "cogito ergo sum" (Descartes, 1637/1994) and that "*concepts* are the building block of thought" (Solomon, Medin, & Lynch, 1999, p. 99). If these statements are true, then concepts are a prerequisite for our existence. A concept is the knowledge or beliefs a person has about members (objects) of a real-world category. Since the ways in which we interact with the world influence what information we store about it (Solomon et al., 1999), people operating in the same environment, but with different tasks and goals, will often form concepts with dissimilar content.

There is a dynamic relationship between concepts and the environment. Concepts determine how people choose to act in the environment at the same time as they are shaped by the interaction that emerges. A frequently debated question is whether all concepts are represented qualitatively equal or if there are qualitatively different modes of representations. One theme in the literature highlights the importance of multiple representation systems that is adaptively activated depending on the task at hand (e.g., Ashby, Alfonso-Reese, Turken, & Waldron, 1998; Juslin, Olsson, & Olsson, 2003; Juslin, Jones, Olsson, & Winman, 2003; Johansen & Palmeri, 2002). On the other hand, others have argued that different tasks will direct the attention to different aspects of the stimulus of the environment and that this alone can explain the results that have been attributed to uses of different representations (e.g., Nosfsky & Johansen, 2000; Pothos, in press).

A. B. Markman, B. H. Ross, and colleagues have provided recent contributions to the literature on how concepts are formed, stored, and used. In a range of studies they trained participants in the same environment using either an *inference learning* paradigm or a *categorization learning* paradigm. The aim has been to study how these different learning paradigms shape the concepts (for a review see Markman & Ross, 2003). Among the main findings is that inference learners are much more sensitive to within-category correlations than categorization learners (Chin-Parker & Ross, 2002) and that while categorization learners appear to store information about diagnostic features or exemplars, inference learners store the prototypical value on each feature dimension (Anderson, Ross, & Chin-Parker, 2002).

The idea that there is a fundamental difference in how inference learning and categorization learning shape our concepts has important implications. The majority of research on how concepts are formed and used has focused on categorization tasks (Solomon et al., 1999). If concepts are formed in a qualitatively different way when people learn through inference, then the contemporary understanding of how concepts are formed and used must be revised.

The present study address the question of whether the differences found between inference learning and categorization learning are likely to be due to a representational shift or a shift in attention. The inference learning task used is hypothesized to focus the attention of the participants differently than the inference learning tasks typically used in the literature.

Inference vs. Categorization

A categorization is defined as a judgment where one or more cues of a stimulus are known and the task is to predict the category label. An inference is defined as a judgment where the category label together with one or several cues of a stimulus is known and the task is to predict the value on an unknown target cue (Markman & Ross, 2003).

The most investigated category structure in inference vs. categorization experiments is the linearly separable (i.e., a category structure in which an additive evidence rule can correctly classify all exemplars) Family Resemblance Category structure (FRCS) (e.g., Anderson et al., 2002; Chin-Parker & Ross, 2002, 2004; Yamauchi & Markman, 1998, 2000). For an example of an FCRS, see Table 1 (the FRCS

used in Yamauchi & Markman, 1998). A FRCS includes two categories (A and B) and the exemplars have a set of binary cues. All members of a category share all but one cue-value with the prototype of its category. The two prototypes typically act as critical exemplars in the test-phase, that is, exemplars for which the difference between the judgments made by inference learners and the judgments made by categorization learners is predicted to be especially large. These critical exemplars are withheld in the learning-phase.

Table 1: A family-resemblance category structure with two categories, eight four-featured Learning Exemplars (LE) and one prototype (Pr) for each category.

		Cat A				Cat B		
Pr	0	0	0	0	1	1	1	1
LE	0	0	0	*1*	1	1	1	*0*
	0	0	*1*	0	1	1	*0*	1
	0	*1*	0	0	1	*0*	1	1
	1	0	0	0	*0*	1	1	1

Categorization Learners (CL) are shown an exemplar and they are asked to predict the category label (i.e., A or B). For perfect performance in the FRCS the participants have to focus on at least three cues. If they focus on only one cue, for example judging A each time cue X has value 0, they will make incorrect judgments in 20% of the cases (Markman & Ross, 2003).

Inference Learners (IL) are shown all but one cue for an exemplar together with the category label. The task is to predict the value on the missing cue (the target cue). For perfect performance, focus on the other three cues is needed. To see this, consider the inference judgment on the first cue dimension on exemplar [1,0,0,0]. To know that the correct answer is 1, the participant has to take into consideration that all other cues have value 0. As for CL, if IL rely on only the relationship between target cue and category label they will be wrong in 20% of the judgments. Thus, the relationship between each cue and category label is probabilistic in both conditions.

Shift in Representation or Attention?

A potential limitation of the designs typically used to investigate inference learning and category learning is that the participants never make inferences for exception features (the italicized values in Table 1) in the learning-phase. The reason for this is that inference judgments on exception features are argued to be analogous to a categorization judgment of the prototype (Yamauchi & Markman, 1998). Since the two prototypes are to act as critical exemplars in the test-phase, inference judgments of exception features is excluded in the learning-phase. Unfortunately, this has the effect that if IL learns to ignore the presented exemplar and only choose on the basis of the category label, the task becomes deterministic. If they choose value 1 on cue X every

time category label = B, they will make no errors. Thus, the categorization task is probabilistic when focusing on only one cue, while there is a deterministic relationship between cue and category label in the inference task if the participants learn to ignore the exemplars.

An alternative to the claim that the inference task and the categorization task "has been demonstrated to yield very different category representations" (Markman & Ross, 2003, p. 592) is that the differences between the two are due to attention being focused differently in the two tasks. If we instead replace the standard inference learning task with a probabilistic one that includes inference judgments of exception features, we should be able to move the participants' attention toward the cues of the exemplars. If how concepts are formed is a question of where the attention is focused this manipulation should change the way IL make judgments in the test-phase. If qualitatively different representations are formed for IL and CL the differences found by Markman, Ross and colleagues should persist also after this manipulation.

The Present study

This study focuses on the generalizability of three of the main findings in the literature on inference vs. categorizations. First, IL learns linearly separable category structures more easily than CL (Yamauchi & Markman, 1998; Anderson et al., 2002). For example, in Yamauchi and Markman (1998) the learning-phase ended when the participant proportion of correct judgments in the last 24 judgments reach .9. That the criterion is reached faster by IL than by CL is argued to be an effect of that IL learns linearly separable category structures more easily. However, remember that if IL ignore the exemplars and judge on the basis of target cue and category label they will be able to reach this criterion by only focusing on one cue. CL on the other hand, have to incorporate knowledge about three cues to reach this criterion. Thus, the difference between the two conditions could be due to the fact that one of the learning tasks is deterministic while the other is probabilistic. In this study, the inference learning-phase task is made probabilistic by including judgments of exception features. As a result, to reach perfect performance, also IL have to take several cues under consideration in every judgment. Despite of this change, will IL still learn the linearly separable category structure more easily?

Second, while CL appear to store information about diagnostic features or exemplars, IL appear store the prototypical value on each feature dimension (Yamauchi & Markman, 1998; Anderson et al., 2002). In Anderson et al. (2002) participants in both conditions carried out a one-cue categorization task in the test-phase. A stimulus with one cue was presented and the participants were to categorize it as A or B. In both Experiment 1 and 2 IL categorized the one-cue exemplar as belonging to the category for which that cue-value was prototypical more often than CL did (percentage of times the category for which the cue was

prototypical was chosen; Exp. 1: IL = .84 and CL=.62; Exp. 2: IL = .92 and CL=.70). This was taken as support for the hypothesis that IL stores the prototypical values. However, an alternative hypothesis is that since the inference learning task is deterministic, IL give more consistent judgments in the test-phase. Since the task in the present experiment includes exception cue inferences, and thereby makes the inference learning task probabilistic, it provides a stronger test of the claim that IL primarily store the prototypical values.

Third, IL appears to learn within-category correlations while CL does not (Chin-Parker & Ross, 2002). Within-category correlations are correlations that do not add to the predictiveness of the features (Chin-Parker & Ross, 2002). In other words, the within-category correlation is uninformative in a categorization task but it is informative in an inference task. That IL are much more sensitive to within-category correlations is interesting since if it is true, it accentuates the importance of using a range of learning paradigms in research on concept formation and usage. People are obviously capable of learning within-category correlations. That they have trouble learning them through categorization indicates that they learn different aspects of an environment through different interactions.

Method

Participants Forty undergraduate students (17 men and 23 women) in the age of 18 to 46 (average age = 24) participated. The participants were compensated with one or two movie tickets depending on their performance in the learning-phase.

Materials The experiment was carried out on a PC-compatible computer. The stimuli used were descriptions of 30 six-featured companies (see Table 2). The companies were all different combinations of the same six binary cues (se Table 3). The cues where chosen to be as neutral as possible.

Design and procedure The category structure includes two linearly separable categories, A and B (see Table 2). There are six binary cue dimensions (C1-6, C1 is the leftmost cue). Among the twenty-four exemplars two are prototypes (prototype of category A = [000000] and the prototype for Category B =[111111]). The prototypes are maximally separated (i.e., they share no features). In the learning-phase participants are presented with the fourteen exemplars on the top of Table 2. The exemplars vary in the frequency of presentation and there are two times as many Category B exemplars as there are Category A exemplars. The cue-validity of the cues vary, C1-C2 have a cue-validity of .9, C3-C4 have a cue validity of .8, and C5-C6 have a cue validity of .7.

In the test-phase task ten old exemplars from the learning-phase (marked with * in Table 2) were presented together with ten new exemplars. Category B has the highest base-rate also in the test-phase.

Note that among the learning-phase exemplars there is a perfect within-category correlation between C5 and C6.

When making inferences on C5 and C6 knowledge about this correlation is crucial. In determining whether an exemplar belongs to A or B, however, knowledge about the correlation is of little value.

There were two conditions, one Categorization-Learning condition (CLC) and one Inference-Learning condition (ILC). The experiment included a learning-phase and a test-phase. In the test-phase the participants were introduced to a computerized trainee-program that would teach them about stock development on a fictive market. In the CLC the participants were presented with a six-featured company and were to judge if the stock of that company had increased or decreased in value during the last year. During three seconds, directly following the judgment, they received feedback concerning the actual stock-value development (the feedback was written in red if the answer had been wrong and in green if it had been correct). Participants made 180 judgments in the learning-phase (the 30 exemplars at the top of Table 2 presented in six consecutive blocks).

Table 2: The category structure used in the experiment. Top: 14 unique exemplars presented in the learning-phase (*f* = times presented in one learning block), divided into two categories (A and B). Exemplars marked with a star are also presented in the test-phase. Bottom: Exemplars presented in the test-phase only.

Presented in training			
Category A		Category B	
Exemplar	*f*	Exemplar	*f*
000000*	2	111111*	2
000011*	3	110111*	2
001000	1	110011*	2
010000*	1	101111*	2
000100	1	011111*	2
001100*	1	111011	4
100000	1	111100*	6
Included in test-phase			
Category A		Category B	
001011		001101	
100011		101100	
111000		110100	
110010		110001	
		010101	
		000111	

A company presented in the ILC included information about five cue-values and category label. The task was to judge which of the two values on the missing binary cue dimension that belonged to the company. Feedback was given as in the CLC. They made one decision for each cue dimension for each company (exemplars in Table 2), thus, they also made a total of 180 judgments.

Exemplar presentation order, which cue in Table 3 that belonged to which cue dimension in Table 2, and which cue value that more often belonged to companies with stocks that increased in value were randomized across all partici-

pants. Due to difference in base-rates of exemplars belonging to category A and B (Table 2), Category A included the companies with positive stock-value development for half the participants (ten in each condition) and the companies with a negative one for the other half.

Table 3: Descriptions of the 6 binary cues.

Descriptions
1) Listed at the NKB/IPEK stock exchange
2) Less/more than 1000 employees
3) Give money to charity/sponsor sports team
4) Active in specific region/whole country
5) Is in the service/industrial sector
6) Primarily export-based/import-based

The test-phase, identical for the two conditions, consisted of four probability judgment tasks[1]. In the first test-phase task, one cue value was presented and the task was to judge the probability that the stock of a company with that cue had increased/decreased in value during the last year (p(cat.│cue)). The participants made 12 judgments (both p(A) and p(B) for the two cue values on C1, C3, and C5 in Table 2). The second test-phase task was the complete opposite. The participants were to judge the probability of cue-value given category (p(cue│cat.)) (12 judgments). The third test-phase task was an inference judgment task. Companies were presented as in the inference learning condition and the participants judged the probability that it had a particular value on C6. In the third and fourth test-phase tasks the participants made two probability judgments for each of the ten old and the ten new exemplars (see Table 2) (40 judgments). The forth task was similar to the third but here the participants were only presented with the five cues and not prior stock development (40 judgments).

In all four tasks in the test-phase the participants were told to give their answer in percent and even up to 0, 10,…100. The entire procedure took about 45 min - 1 ½ hour.

Results and Discussion

Did IL learn more easily? In the last 60 judgments the participants in the CLC made significantly more correct judgments than did the participants in the ILC (proportion of correct judgments among the last 60 judgments; CLC = .93 and ILC = .69; t_{38}=7.401, p<.001). In contrast to earlier

---

[1] When the number of judgments for each exemplar in the test-phase is low, probabilities are a more sensitive measurement than forced choices. For example, if the participants makes two judgments for one exemplar the proportion of times they choose A rather than B can only be 1.0, .5 or .0, the mean probability judgment, on the other hand, can take any value between 1.0 and 0. Usually in studies on inference vs. categorization the task is to choose category or cue value rather than judging the probability of cue value or category. It has been shown, however, that there is a high correlation between response proportions and probability judgments (Peterson & Beach, 1967).

results, this linearly separable category structure was easier to learn in the CLC. Thus, when both learning tasks are probabilistic it appears as if it is harder to learn from category to cue than from cue to category (for similar findings see Nilsson & Björkman, 1982).
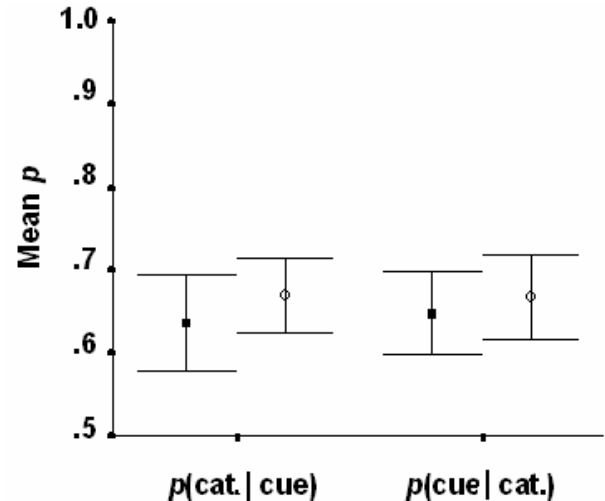

Figure 1. Mean probability judgments with 95% confidence intervals for the probability of category given cue (p(cat.│cue)) and the probability of cue given category (p(cue│cat.)) for both conditions (filled squares=ILC and transparent circles=CLC).

Did IL rely more on the prototypical values? Figure 1 shows both the mean probability judgments that a company with cue-value X belongs to the category for which X is the prototypical cue value (p(cat.│cue)) and the mean probability judgments that a company that belongs to category Z have the feature-value that is prototypical for Category Z (p(cue│cat.)).

In both tasks (p(cat.│cue) and p(cue│cat.)) there is a slight tendency, though far from significant, for the CLC participants to rely relatively more on the prototypical cue-value (p(cat.│cue): t_{38}=.939, p=.35; p(cue│cat.)= t_{38}=.597, p=.55). Note that the tendency is in the opposite direction compared to the trend in Anderson et al. (2000). A power analysis of the one-cue categorization task (p(cat.│cue)) based on the standard deviation taken from the data in this experiment (s=.11) showed power of finding a difference of half the magnitude of the difference in Anderson et al. (2002)[2] in our study was more than 99% (δ=4.43). Thus, by changing the characteristics of the task and the environment, but holding the mode in with inference decisions was made in the learning-phase constant, the relatively higher preference for the prototypical values for IL vanished.

---

[2] In Anderson et al. (2002) the difference in proportion of times the participants choose the category for which the cue was prototypical between IL and CL was .22. Given the high correlation between response proportions probability judgments (Peterson & Beach, 1967), half the difference in Anderson et al. (2002) is assumed to be equal to a difference in probability judgments of .11.
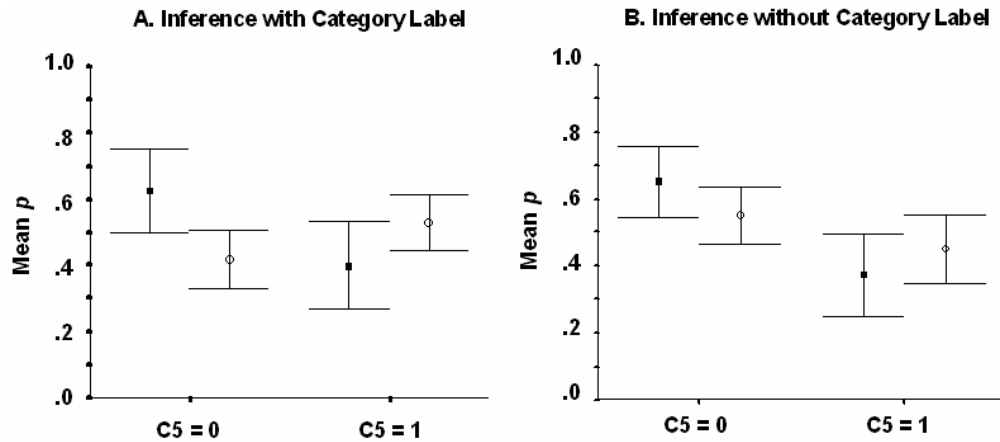
Figure 2. Mean probability judgment, with 95% confidence intervals, that C6=0 (the data from the probability judgments that C6=1 was re-coded in the form of a probability judgment that C6=0) when C5 is either 0 or 1 for both conditions (filled squares=ILC and transparent circles=CLC).

It could be argued that this is due to IL having less knowledge, and if they were trained more they would rely more on the prototypical value. There are two aspects of the data that speaks against such a claim. First, if the participants had made random judgments the mean judgment would end up around .5. As can be seen in Figure 1 the confidence intervals for IL in both test-phase tasks are clearly separated from .5, thus indicating that IL were not ignorant. Second, it should also be noted that for IL the mean judgment of p(cat. | cue) for C1 (p(cat. | cue) =.71) was higher than for C3 (p(cat. | cue) =.68) and the mean judgment of p(cat. | cue) for C3 was higher than for C5 (p(cat. | cue) =.52). That they had learned the correct "hierarchy" of cue validities further indicates that although the participants in the ILC had rather good knowledge of the environment they still did not rely on the prototypical value. A repeated measures ANOVA showed that the difference in p(cat. | cue) between cues was significant in the ILC ($F_{2,38}$=8.804; MSE=.023; p=.001).

Thus, IL participants did not appear to rely relatively more on the prototypical values. Rather than being an effect of qualitatively different cognitive representations the finding that IL focus more on prototypical value could be, at least in part, an artifact of the deterministic learning task that is typically used.

Were the categorization learners incapable of learning the within-category correlations? Figure 2 shows the results from the inference judgments on C6 in the form of probability judgments for new exemplars (i.e., exemplars not seen in the learning-phase), with (A) and without (B) category label. In order to facilitate presentation the data from the probability judgments that C6=1 was re-coded in the form of a probability judgment that C6=0. Figure 2 illustrates to what extent the participants made judgments consistent with the within-category correlation between C5 and C6 that existed in the training-phase. To be consistent with the within-category correlation they should judge the prob-

ability that C6 = 0 to be above .5 when C5=0 and to below .5 when C5=1.

As in previous findings, the participants in the ILC made inference judgments consistent with the within category correlation both when category label was known and when it was unknown. Also consistent with previous finding was that when category label was presented (Figure 2A), the participants in the CLC made inference judgments that were inconsistent with the within-category correlation. However, when category label was not presented (Figure 2B), participants in the CLC made judgments that were consistent with the within-category correlation. Thus, participants in the CLC appeared to use different judgment strategies in the two tasks. This was confirmed by an ANOVA with Value on C5 (1 or 0) and Test-phase inference task (with and without category label) as independent variables and Probability judgment as dependent variable that showed that neither of the two main effects were significant for CLC (Value on C5: $F_{1,19}$=.008, MSE=.089, p=.93; Test-phase inference task: $F_{1,19}$=1.295, MSE=.012, p=.269). However, the interaction effect was significant ($F_{1,19}$=6.948, MSE=.033, p=.02).

Concluding Remarks

It has been argued that inference learning and categorization learning lead to qualitatively different cognitive representations (Markman & Ross, 2003). Phenomenons that have been attributed to this difference are that IL learns linearly separable category structures more easily than CL (Anderson et al., 2002) and that IL stores the prototypical values of the objects of a category (Yamauchi & Markman, 1998). In this study it was shown that these findings do not generalize to a design where both learning tasks have a probabilistic relation between cue and category label. The results of the present study suggest that rather than being due to qualitatively different cognitive representations, the differences that previously have been found in judgments made by IL

and CL were due to attention being focused differently in the two tasks. Our results do not, however, imply that inference learning never leads to reliance on other cognitive representations or memory systems than categorization learning. Recent research shows that a full understanding of category learning must go beyond the explanations provided by traditional categorization models and investigate the relationship between category learning and the major memory systems identified in the literature (Ashby & O'Brien, 2005). In a probabilistic classification task, for example, it seems that people rely on procedural knowledge early in training and on declarative memory later in training (Knowlton, Squire, & Gluck, 1994). The process switching during learning has been confirmed in neuroimaging studies showing changes in the neural networks involved early and late in training (e.g., Poldrack, Prabhakaran, Seger, & Gabrieli, 1999).

Regarding the finding that CL are insensitive to within-category correlation (Chin-Parker & Ross, 2002) the results were more ambiguous. The inference judgments made by the CL were inconsistent with the within-category correlation when category label was provided. However, when less information was provided (when category label was absent) CL made inference judgments consistent with the within-category correlation. Interestingly, inference without category label has as far as the authors know never been tested before. Thus, the results suggest that the claim that CL are insensitive to within-category correlations might be premature. However, the question of when CL uses knowledge of within-category correlations has to be investigated further.

This study show that several of the findings previously attributed to differences in cognitive representation can be better explained by how attention is focused in the different learning tasks. The results indicate that the differences found might just as well be explained by a shift in attention as by a shift in representation.

References

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

Ashby, F. G., Alfonso-Reese, L. A., Turken, A. U., & Waldron, E. M. (1998). A neuropsychological theory of multiple systems in category learning. *Psychological Review, 105*, 442-481.

Ashby, F. G., & O'Brien, J. B. (2005). Category learning and multiple memory systems. *Trends in Cognitive Sciences, 9,* 83-89.

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 categorization learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 30*, 216-226.

Descartes, R. (1994). *Discours de la méthode.* Niotre Dame, University of Notre Dame Press (original work published 1637).

Johansen, M. K., & Palmeri, T. J. (2002). Are there representational shifts during category learning? *Cognitive Psychology, 45*, 482-553.

Juslin, P., Jones, S., Olsson, H., & Winman, A. (2003). Cue abstraction and exemplar memory in categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 29,* 924-941.

Juslin, P., Olsson, H., & Olsson, A-C. (2003). Exemplar effects in categorization and multiple-cue judgment. *Journal of Experimental Psychology: General, 132*, 133-156.

Knowlton, B. J., Squire, L. R., & Gluck, M. A. (1994). Probabilistic classification in amnesia. *Learning and Memory, 1,* 106-120.

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

Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning, *Psychological Review, 85*, 207-238.

Nilsson, R., & Björkman, M. (1982). Prediction and diagnosis: Task variable revealed. *Scandinavian Journal of Psychology, 23*, 253-262.

Nosofsky, R. M., & Johansen M. K. (2000). Exemplar-based accounts of "multiple-system" phenomena in perceptual categorization. *Psychonomic Bulletin & Review, 7*, 375-402.

Peterson, C. R., & Beach, L. R. (1967). Man as an intuitive statistician. *Psychological Bulletin, 68*, 29-46.

Poldrack, R. A., Prabhakaran, V., Seger, C. A., & Gabrieli, J. D. E. (1999). Striatal activation during acquisition of a cognitive skill. *Neuropsychology, 13*, 564-574.

Pothos, E. M. (in press). The rules versus similarity distinction. *Behavioral and Brain Sciences.*

Solomon, K. O., Medin, D. L., & Lynch, E. (1999). Concepts do more than categorize. *Trends in Cognitive Sciences, 3*, 99-105.

Yamauchi, T., & Markman, A. B. (1998). Category learning by inference and classification. *Journal of Memory and Language, 39*, 124-148.

Yamauchi, T., & Markman, A. B. (2000). Inference using categories. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26*, 776-795.