

From Prototypes to Exemplars: Representational Shifts in a Probability Judgment Task

Håkan Nilsson (hakan.nilsson@psy.umu.se)

Department of Psychology, Umeå University
SE-901 87, Umeå, Sweden

Peter Juslin (peter.juslin@psy.umu.se)

Department of Psychology, Umeå University
SE-901 87, Umeå, Sweden

Henrik Olsson (henrik.olsson@psyk.uu.se)

Department of Psychology, Uppsala University
SE-741 42, Uppsala, Sweden

Abstract

What information is used to make subjective probability judgments? In this study we test a hypothesis, grounded both in research on categorization and developmental psychology, proposing that when first confronted with an environment people create prototypes and as a function of learning they start to store concrete exemplars. The hypothesis implies that this representational shift will appear later in an environment with more complex stimuli. Therefore the hypothesis was tested with both a standard stimulus sample and a complex stimulus sample. The results indicate that more people rely on prototypes at an early stage in learning than later on, and that the shift towards more widespread exemplar memory reliance appears later in an environment with more complex stimuli.

Introduction

People constantly find themselves in situations where they have to choose one of several alternatives as, for example, when a physician has to diagnose a patient. The physician considers possible diseases and decides which disease is the most likely. This involves assessment of the *subjective probability* that the symptoms of the patient are symptoms of a particular disease. Previous cognitive theories have emphasized that probability judgments derive from relative frequencies (e.g., Gigerenzer, Hoffrage, & Kleinbölting 1991; Juslin, 1994) or representativeness, commonly conceived of as the similarity between a probe and the category prototype (Kahneman, Slovic, & Lichenstein, 1982). More recent studies have, however, favored *exemplar memory* as the basis for subjective probability (Juslin & Persson, 2002; Juslin, Nilsson, & Olsson, 2001; Sieck & Yates, 2001; see also Dougherty, Gettys, & Ogden, 1999).

Since the eighties one of the most prominent theories of categorization has been *exemplar theory* (Medin & Schaffer; 1978; Nosofsky & Johansen, 2000). The theory states that when categorizing a patient, the diagnosis is determined by the similarity of the new patient to memory traces of the individual exemplars belonging to different disease categories. Minda and Smith (2001), however, argue that the out-

standing success of exemplar theory, in part, derives from the specific category structures used in most previous studies. They suggest that the reliance on exemplars is preceded by a stage involving category prototypes. The basic assumption in their argument is that a “novice” bases the judgments on the similarity to a category prototype, while an “expert” relies on similarity to individual category exemplars. If there is a difference between novices and pros, then the proposed shift should appear later in an environment with more complex stimuli.

A shift from prototypes to exemplars is observed also in research on early word acquisition. It is proposed that when learning a word the child initially creates a prototypical representation that has the central features of the objects in the natural category. The more similar an object is to the prototype for a word the more likely the child is to use that word to name the object. New words (exemplars) for objects within a category are learned and defined in relation to their similarities and dissimilarities to the prototype (Barret, 1986).

If this representational shift reflects a primary aspect of the biological function of the brain, it is reasonable to hypothesize that it is also a central aspect of the way adults acquire knowledge. The goal of this study is to investigate whether such a shift exists and if it appears later in an environment with more complex stimuli. That is with more features and more distinctive feature combination. To do this, two conditions with the same category structure but with different stimuli constitution (standard and complex) was tested. We thus ask: Is there a representational shift from reliance on the representativeness heuristic (i.e., prototype similarity) to exemplar memory in probability judgment as participants develop from “novices” to “pros” in the task?

Exemplar Theory

Both the exemplar and the prototype theory emanate from the categorization literature and assume that responses are based on the similarity between a new probe and category representations (exemplars or prototypes) stored in memory. Exemplar theories suggest that categories are composed of

exemplars of events or objects that we have encountered. Not all events or objects are stored and the ones stored are not necessarily complete representations of events or objects. The crucial point is that the representations are of concrete events or items and as a function of exposure to a specific environment the number of stored exemplars will increase. In contrast to most applications of exemplar and prototype models, in this article these are used to model subjective probability judgments rather than categorizations.

Consider the following; a physician is assigned to judge the probability that a patient has got one of two diseases. Exemplar theory states that to do this he or she retrieves exemplars of previous patients that had disease *A* and of previous patients that had disease *B*. He or she then compares the new patient to these exemplars and bases his or her judgment upon the similarity to the exemplars of disease *A* and *B*. The probability judgment will reflect the *similarity* to the exemplars of disease *A* relative to the similarity to the exemplars of disease *B*, but it will also reflect the *frequency* of previous patient that have had either of the two diseases.

The exemplar model used in this paper is a simplified version of *PROBEX* (i.e., PRObabilities from EXamplares; Juslin & Persson, 2002), which in turn modifies the original *context model* (Medin & Shaffer, 1978) for application to subjective probability judgment. The version used here amounts to using the context model to predict subjective probability judgments rather than response proportions. The probability that a new item *t* belongs to category *A* is given by Equation 1.

$$p(A) = \frac{\sum_i S(t|x_i)c(x_i)}{\sum_i S(t|x_i)}, \quad (1)$$

Here x_i are the retrieved exemplars ($i = 1 \dots I$), $c(x_i)$ is determining which category the retrieved exemplar belongs to ($c(x_i)=1$ if x_i belongs to Category *A* and $c(x_i)=0$ if x_i belongs to Category *B*), and $S(t|x_i)$ is the similarity between the item and the retrieved exemplar. The similarity is computed by the multiplicative similarity rule of the context model (Medin et al., 1978):

$$S(t, x) = \prod_{j=1}^D d_j, \quad d_j = \begin{cases} 1 & \text{if } t_{.j} = x_{.j} \\ s & \text{if } t_{.j} \neq x_{.j} \end{cases}, \quad (2)$$

where D is the number of features (in this study $D=4$), d_j is 1 if the values on feature j match and s if they mismatch. The similarity value s is a free parameter that can take any value between 0-1 and reflects the impact of a certain mismatching feature. Thus, the physician would compare new patient (*t*) with old patients (x_i) and judge the probability of disease *A* on the basis of the similarity to the retrieved exemplars.

Prototype Theory

Prototype theory also states that the response is based on the similarity to a category member. The difference is that while the categories in exemplar theory contain several members, prototype theory suggests that each category is represented by one member (the prototype). The prototype is an abstraction that reflects the most common features among members of a category in the environment (or at least of the ones en-

countered; e.g., Smith & Minda, 1998). According to the representativeness heuristic the physician would compare a new patient with the prototype for disease *A* patients with the prototype for disease *B* patients and base the probability judgment on the similarity to these two (Kahneman, Slovic, & Tversky, 1982). The probability for an item *t* to belong to category *A* is:

$$p(A) = \frac{S(t|P_A)}{S(t|P_A) + S(t|P_B)}, \quad (3)$$

where P is the prototype, S is the similarity of *t* to P (S is derived in accordance to Eq. 2). Thus with prototype theory $p(A)$ does not reflect frequency in the environment, only the similarity to the most central features within a category.

The Representational Shift

Why would there be a representational shift from prototype to exemplar reliance in a subjective probability judgment task? In recent years results have emerged in categorization research indicating a shift from abstract representation, either as rules (Johansen & Palmeri, 2002) or prototypes (Smith & Minda, 1998), to exemplar representation. Given the close relationship between the two fields, if the shift exists in categorization it should appear also in subjective probability judgment. A theoretical suggestion why such a shift should appear is that humans tend to assume natural categories to be linearly separable and decision boundaries to be linear (Smith & Minda, 1998). In such an environment it is not efficient to memorize separate exemplars at an early stage. Instead it is preferable to abstract rules or prototypes that catch central aspects of the category. It is not until this category-knowledge is stored that it is possible to easily discriminate the features within each exemplar.

If there is a shift the initial representation could either be rules or prototypes. A suggestion that the initial representation is in the form of prototypes is provided in the field of early lexical development. It is proposed that when learning the meaning of a word, e.g. "bird", the child will pair the word with an object that has got the important feature of the "bird-category", a bird prototype. As a function of maturity the child will first expand its usage of the word to objects in the category that share many features with the prototype and then on to objects of the category that share fewer features with the prototype. Exemplars take over objects previously represented by the prototype which decreases in importance (Barret, 1986). This theory has been prosperous in explaining a wide range of frequently observed phenomena in research on lexical development (see Barret, 1996; Meints, Plunkett, & Harris, 1999; Southgate & Meints, 2000).

Complexity

Minda and Smith (2001) criticized the designs normally used in research on categorization, claiming that exemplar theory is given such an advantage that a prototype-phase never gets a chance to be discovered. They argue that the category structures normally have featured categories with a

few low-dimensional stimuli, and that this relative simplicity gives an incitement for exemplar storage.

It is no bold proposition that real-world categories are normally much more complex than categories in experimental settings (the control of variables is no less than the *raison d'être* of experiments), and if the success of the exemplar theory is due to simplicity in experimental categories then that would be an important finding. In real-world categories the same feature may have a wide range of appearances, e.g. all humans have eyes but they differ in color and size. Yamauchi and Markman (2000) have shown that if each feature in a standard categorization task (as in e.g. Minda & Smith 2001; or Medin & Shaffer, 1978) is given multiple levels the difficulties of learning increases severely.

In order to examine the effect of an increase in stimulus complexity we propose a combination of the arguments by Minda and Smith (2001) and the findings of Yamauchi and Markman (2000). As mentioned, the previous authors argue that a more complex category structure with more features and more exemplars will give the prototype-phase a more fair chance of being discovered. By increasing the complexity in the same manner as in Yamauchi and Markman (2000), by introducing multiple feature-levels, this is achieved. In the eyes of the participant the number of features will increase which will increase the number of distinctive exemplars.

Category Structure

The core hypothesis that is tested in this paper is whether there is a representational shift from reliance on prototypes to reliance on exemplar memory in a probability judgment task. Therefore a category structure designed to give both quantitatively and qualitatively different predictions by the two models was created. The category structure is shown in Table 1, and it contains a total of 60 exemplars equally divided across two categories. There are 12 distinctive feature combinations and each exemplar has got four binary features.

Table 1: The 12 distinctive exemplars and how frequently they appear in each category.

Category A		Category B	
Exemplar	Freq.	Exemplar	Freq.
1 1 1 0	1	1 1 1 0	5
1 0 0 0	5	1 0 0 0	1
1 1 0 1	1	0 0 1 0	1
1 0 1 1	1	0 1 0 0	1
0 1 1 1	6	0 0 0 1	6
1 0 0 1	1	1 0 0 1	1
0 1 1 0	1	0 1 1 0	1
1 1 1 1	14	0 0 0 0	14

The two categories were created around two well separated prototypes (1111 for category A and 0000 for category B). To enable a qualitative differentiation between the two models there are two critical exemplars (1000 and 1110).

These are each very similar to one prototype but appears more frequently in the other category. Thus, the prototype model will predict the judgment for these exemplars to be based on the similarity to the prototype whereas PROBEX predicts that the judgment will be based on the similarity and frequency of exemplars in the two categories. The a priori predictions are presented in Figure 1. The interaction in Figure 1 is valid as long as $s < .3$, with $s > .3$ the prediction of the prototype model for feature combination 1110 will always be higher than the prediction by PROBEX, and vice versa for feature combination 1000.

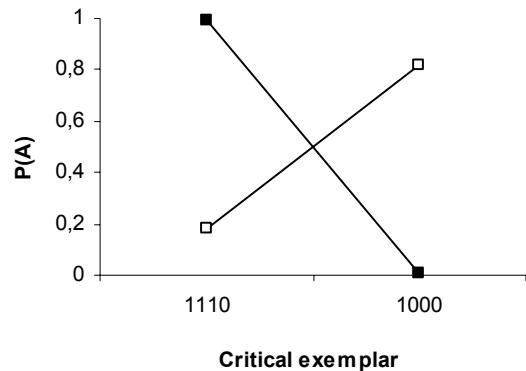


Figure 1: Predictions by PROBEX (white squares) and the prototype model (black squares) for the two critical exemplars (the predictions are calculated with $s=0$).

To evaluate the claim that in a more complex environment people will rely on prototypes for a longer period, there was one *standard stimulus condition* and one *complex stimulus condition*. In both conditions the same category structure is used (Table 1). In the standard condition the stimuli were constituted as they traditionally have been in standard categorization tasks, that is few in number and with a low amount of features. In the complex condition stimuli was presented as in Yamauchi and Markman (2000), that is with multiple levels (six) of each feature, three corresponding to a 1 in Table 1 and three corresponding to a 0 in Table 1. This makes it possible to create 81 ($3^3 \times 3^3$) unique stimuli equivalent to each of the 12 distinctive exemplars in Table 1. In other words, even if the complex condition is based on the same category structure as the standard condition the environment is more complex since the same stimuli never reappear.

Method

Participants

Thirty undergraduate students (15 men & 15 women) in the age of 19 to 28 (average age = 23) participated. The participants were paid between 50-100 SEK depending on how well their probability judgments were calibrated.

Materials and Procedure

The experiment was carried out on a PC-compatible computer. The stimuli used were subspecies of the sometimes

poisonous (but fictive) Deathbug (see Figure 2). There were 12 distinctive four-featured subspecies in the standard condition (see Table 1), the features were “leg length and colour” (long light grey vs. short dark grey), “colour and pattern on posterior body” (no pattern on a green body vs. white dots on brown body), “nose length and colour” (long grey vs. short blue), and “dots or no dots on anterior body” (red dots on beige surface vs. no dots on beige surface). The total number of stimuli in Table 1 (60) was presented four times (creating four blocks).

In the complex condition there were 240 distinctive subspecies (equivalent to the 60 exemplars in Table 1 presented four times). The core features (leg length, nose length, dots or no dots, and posterior body pattern) were identical to those in the standard condition but the color of the body parts had three levels of strength (e.g. the short nose could be light blue, standard blue or dark blue). To further increase the uniqueness of each presented bug they were haphazardly spray painted on the back and each presentation was rotated (there were 62 possible positions, with a maximum rotation of 90 degrees in each direction).

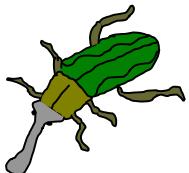


Figure 2: An example of one four-featured subspecies of the Deathbug, as presented to the participants.

Design and procedure

The design was a 2*2 mixed design with probability judgments as dependent variable and stimulus complexity (two levels, between-subjects) and training (four blocks, within-subjects) as independent variables.

The participants were presented with subspecies as the one in Figure 2. Initially they were to judge whether this particular subspecies was dangerous or not (category A or B in Table 1). When a decision had been made the participants judged, on a scale from 50-100%, how confident they were that it was dangerous/not dangerous (this judgment depended on the previous answer). Following each of these judgments the participants received feedback on the actual dangerousness of the previously judged subspecies (this feedback was given when the subspecies was still present on the screen).

In both conditions the participants encountered 240 bugs. The order in which they were presented was randomized across the participants.

Results

Qualitative fit. As mentioned, there were two critical exemplars for which the two models *a priori* provided contradictory predictions (see Figure 1). The mean probability assessment of the critical exemplar 1000 was subtracted from the mean probability assessment of 1110. The differ-

ences across all four blocks are presented in Figure 3. Positive values indicates that the judgments have been based on similarity to prototypes while negative values indicate that the judgments have been made on the basis of similarity to stored exemplars.

In both conditions, data in Block 1 suggests that the early judgments were based on similarity to prototypes. Data in the later blocks implies that an increasing number of participants base their judgment on similarity to stored exemplars when knowledge increases. In other words, there is an overall trend towards more widespread reliance on exemplars in both conditions. The data in Block 4, however, indicates that the majority of the participants in the complex condition still base their judgments on prototype similarity while the majority of the participants in the standard condition base their judgment on exemplar similarity. Thus, the data in Figure 3 support the hypothesis of a representational shift as a function of learning and that this shift appears later in a more complex environment.

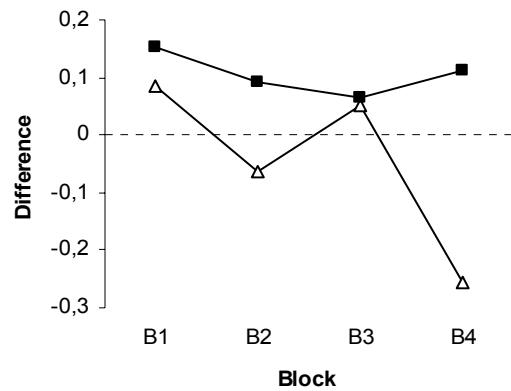


Figure 3: The mean probability assessment for critical exemplar 1000 subtracted from the mean probability assessment for critical exemplar 1110 for both conditions (standard condition = white triangles and the complex condition = black squares) across all four blocks.

Quantitative fit. To examine the quantitative fit of the models they were fitted to the data for all 12 exemplars (remember that also in the complex condition there were only 12 distinctive feature combinations) from each individual participant with Root mean Square Deviation (RMSD) as error function. The models were fitted with one *s*-value for all four features in order to minimize possible problems of over fitting.

The mean of the individual RMSD of PROBEX was subtracted from the mean individual RMSD of the prototype model; this is illustrated in Figure 4.

PROBEX shows better fit in all blocks (except for Block 1 in the complex condition). This is in part at contrast with the findings in Figure 3. However, the overall trend in Figure 4 is equivalent to the trend in Figure 3, hence it indicates that there is a shift towards more extensive exemplar reliance. Data in Figure 4 indicate a mix of strategies among the participants, some relying on prototypes and some on

exemplars, and that relatively more people rely on prototypical representations early on and relatively more people rely on exemplar representation at a later stage.

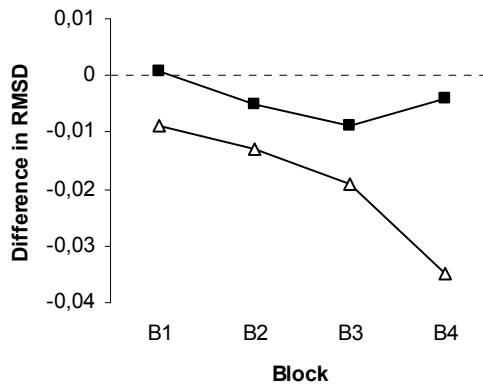


Figure 4: Mean RMSD for PROBEX subtracted from the mean RMSD for the prototype model for both conditions across all blocks, when the models were fitted to individual data. (standard condition = white triangles and the complex condition = black squares) across all for blocks.

Both data in Figure 3 and 4 indicate that if a shift exists, it appears much later in the complex condition than in the standard condition. An ANOVA with Block and Condition as independent variables and RMSD as dependent variable was performed to examine if there were a significant difference between the two conditions. No interaction effect was found ($F_{3,84}=2.015, p=.12$), but as hypothesized both main effects were significant (Block: $F_{3,84}=3.105, p=.03$; Condition: $F_{1,28}=7.502, p=.01$). This further supports the hypothesis that as knowledge increase more and more participants rely on exemplar representation and that the representational shift appears later in an environment with more complex stimuli.

To examine if the findings above was due to increased knowledge, an ANOVA with proportion correct classifications as dependent variable and block and condition as independent variables was performed. Both main effects and the interaction effect were significant (Block: $F_{3,84}=15.95, p<.001$; Condition: $F_{1,28}=3018.3, p<.001$; interaction $F_{3,84}=3.115, p=.03$). Thus, the knowledge had increased from Block 1 to Block 4 and the participants in the standard condition stored information at a higher rate than did those in the complex condition.

In sum, throughout data there is a trend from reliance on prototypes towards reliance on exemplar as a function of learning, indicating that there is a representational shift. Further more this shift appears later in a more complex environment.

Discussion

Does a representational shift from prototypes to exemplars appear in a subjective probability judgment task? There were two prime reasons why we raised this question. First,

in the recent categorization literature there is growing support for a shift with training from abstract knowledge to exemplars (Johansen & Palmeri, 2002; Minda & Smith, 2001). Secondly, a theory of a shift from prototypes to exemplars has been prominent during many years in developmental psychology (Barret, 1996; Southgate & Meints, 2000; Meints et. al., 1999).

If this representational shift exists, what determines when it appears? A theory was tested that hypothesized that the complexity of the environment decides the temporal location of the shift.

Both the qualitative and the quantitative analyses provides support for the hypothesis that a representational shift, from prototypes to exemplars, takes place as a function of learning. Throughout data there is a trend from relative support for the prototype model to relative support for PROBEX as a function of learning.

A difference between the two conditions was found. Participants in the complex condition appeared to rely on prototypical representation for a longer period. Since the knowledge in the complex condition was poorer it is reasonable to believe that the delayed shift was due to this inferiority in stored information.

Therefore, we conclude that the data support the hypothesis of a representational shift from reliance on prototype similarity to reliance on exemplar similarity in a probability judgment task as a person develops from "novice" to "pro". Furthermore, we conclude that, maybe not surprisingly, it takes more training to become pro in a complex environment than in a simple.

Presuming that this shift is a core element in the way the human brain develops an understanding of the world, a sensible (subconscious) strategy is to build up information about what is common to a phenomenon (category). This information can then be used to differentiate between the members within each group. In other words, the exemplars are placed in a context that is necessary to differentiate them from other exemplars of the same category. In the case of the bugs, a prototype containing the central features among dangerous bugs is created and the exemplars are stored in relation to this prototype so that the idiosyncratic features of the exemplars are highlighted. If natural categories mainly are linear separable this is a more efficient strategy than trying to store individual exemplars at an early stage since it separates individual and category features: it is an ecologically plausible tactic.

Even if the results here suggest that shifts from prototype to exemplar representations occur in this task there are some caveats to this interpretation. First, it might be possible for a more complex version of the exemplar model to account for early performance in this task. For example, Nosofsky and Zaki (2002) argued that the exemplar model could account for the data in Minda and Smith (2001) better than a prototype model if the model is supplemented with a response scaling mechanism that already is implicit in the prototype model. To investigate this possibility, however, one must take into consideration the flexibility of the models. The best strategy would be to compare the models with fit indexes that take into consideration both the number of parameters and their functional form (Pitt, Myung, & Zhang,

2002). Second, the results in Johansen and Palmeri (2002) suggest that the shift is from simple rule following to exemplar based processing. The category structures in Johansen and Palmeri, however, were quite different from the one used here and it is therefore difficult to compare the results (i.e., in Johansen & Palmeri, 2002, there were few unique training exemplars and the task involved categorization). However, Johansen and Palmeri faces one problem, and that is to explain the support of an initial prototype-phase found in research on preverbal-children (e.g. Southgate et al., 2000). The findings by Johansen and Palmeri raises the question if there initially is a qualitatively difference in category representation between children and adults.

As discussed earlier, real world categories are more complex than experimental categories. When the complexity of the stimulus was increased so was the support for the shift hypothesizes. But, PROBEX did explain data best in the quantitative analyses also in the complex stimulus condition in at least Block 2-4. One interpretation of this could be that the stimulus was not complex enough. In order to be able to draw more secured conclusions the hypothesis has to be evaluated using even more complex stimulus.

Acknowledgments

The Swedish Research Council supported this research. We are grateful to Morgan Renman for his assistance in conducting the experiment.

References

Barret, M. (1986). Early semantic representation and early word usage. In S. A. Kuczaj II & M. D. Barret (Eds.). *The development of word meaning* (pp. 39-67). New York: Springer-Verlag.

Barret, M. (1996). Early lexical development. In Fletcher, P. and MacWhinney, B. (eds.), *The Handbook of Child Language* (pp. 362-392). Oxford: Blackwell.

Dougherty, M. R. P., Gettys, C. F., & Ogden, E. E. (1999). MINERVA-DM: A memory process model for judgments of likelihood. *Psychological Review*, 106, 180-209.

Gigerenzer, G., Hoffrage, U., & Kleinbölting, H. (1991). Probabilistic mental models: A Brunswikian theory of confidence. *Psychological Review*, 98, 506-528.

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

Juslin, P. (1994). The overconfidence phenomenon as a consequence of informal experimenter-guided selection of almanac items. *Organizational Behaviour and Human Decision Processes*, 57, 226-246.

Juslin, P., Nilsson, H., & Olsson, H. (2001). Where do probability judgments come from? Evidence for similarity-graded probability. In J. Moore & K. Stenning (Eds.), *Proceedings of the Twenty-Third Annual Conference of the Cognitive Science Society* (pp. 471-476). Hillsdale, NJ: Erlbaum.

Juslin, P., & Persson, M. (2002). PROBabilities from EXemplars (PROBEX): A "lazy" algorithm for probabilistic inference from generic knowledge. *Cognitive Science*, 26, 563-607.

Kahneman, D., Slovic, P., & Tversky, A. (Eds.) (1982). *Judgment under uncertainty: Heuristic and biases*. New York: Cambridge University Press.

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

Meints, K., Plunkett, K., & Harris, P. (1999). When does an ostrich become a bird? The role of typicality in early word comprehension. *Developmental Psychology*, 35, 1072-1078.

Minda, P. J., & Smith, J. D. (2001). Prototypes in category learning: The effects of category size, category structure, and stimulus complexity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 27, 775-799.

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

Nosofsky, R. M., & Zaki, S. R. (2002). Exemplar and prototype models revisited: Response strategies, selective attention, and stimulus generalization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 28, 924-940.

Pitt, M. A., Myung, I. J., & Zhang, S. B. (2002). Toward a method of selecting among computational models of cognition. *Psychological Review*, 109, 472-491.

Sieck, W. R., & Yates, J. F. (2001). Overconfidence effects in category learning: A comparison of connectionist and exemplar memory models. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 27, 1003-1021.

Smith, J. D., & Minda, J. P. (1998). Prototype in the mist: The early epochs of category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24, 1411-1436.

Southgate, V., & Meints, K. (2000). Typicality, naming, and category membership in young children. *Cognitive Linguistics*, 11, 5-16.

Yamauchi, T. & Markman, A.B. (2000). Learning categories composed of varying instances: The effect of classification inference, and structural alignment. *Memory and Cognition*, 28, 64-78.