

The Emergence of Perceptual Category Representations During Early Development: A Connectionist Analysis

Paul C. Quinn

Department of Psychology
Washington & Jefferson College
Washington, PA 15301, USA
pquinn@vms.cis.pitt.edu

Mark H. Johnson

MRC Cognitive Development Unit
4 Taviton Street
London, WC1H 0BT, UK
mark@cdu.ucl.ac.uk

Abstract

A number of recent studies on early categorization suggest that young infants form category representations for stimuli at both global and basic levels of exclusiveness (i.e., mammal, cat). A set of computational models designed to analyze the factors responsible for the emergence of these representations are presented. The models (1) simulated the formation of global-level and basic-level representations, (2) yielded a global-to-basic order of category emergence and (3) revealed the formation of two distinct global-level representations -- an initial "self-organizing" perceptual global level and a subsequently "trained" arbitrary (i.e., non-perceptual) global level. Information from the models is used to make a number of testable predictions concerning category development in infants.

Introduction

Investigators interested in early cognitive development have been examining the origins and development of complex category representations during the first two years of life (e.g., Mandler, Bauer, & McDonough, 1991; Mervis, 1987; Quinn, Eimas, & Rosenkrantz, 1993). Empirical efforts have been concerned with (1) the age and means by which individuated representations can be formed for basic-level categories (e.g., cats, chairs) from the same global-level structure (e.g., mammal, furniture), and (2) whether these representations begin to cohere to form global-level representations or whether global-level representations precede basic-level representations.

One set of relevant studies has shown that young infants can form perceptually-based category representations at both basic and global levels of exclusiveness (reviewed in Quinn & Eimas, in press-a). At the basic level, 3- to 4-month-olds familiarized with domestic cats will generalize familiarization to novel cats, but dishabituate to birds, horses, dogs, tigers and even female lions. The data provide evidence that the infants can form a representation for cats that includes novel cats, but excludes exemplars from a variety of related basic-level categories. Behl-Chadha (in press) has extended these findings to human-made artifacts by showing that 3- to 4-month-olds can form individuated representations for chairs and couches each of which excludes instances of the other as well as beds and tables.

At the global level, 3- and 4-month-olds familiarized with instances from a number of mammal categories (e.g., cats, dogs, tigers, rabbits, zebras, elephants) generalized familiarization to novel mammal categories (e.g., deer), but dishabituated to instances of birds, fish and furniture (Behl-Chadha, in press), indicating that the infants can form a

representation of mammals that includes novel mammal categories, but excludes instances of non-mammalian animals (i.e., birds and fish) and human-made artifacts (e.g., furniture). In the same series of experiments, Behl-Chadha obtained evidence that 3- to 4-month-olds can also form a representation for furniture that includes beds, chairs, couches, dressers, and tables, but excludes the mammals mentioned above. The evidence thus suggests that young infants can form global-level representations for at least some natural (i.e., mammals) and artifactual (i.e., furniture) categories.

Of interest is the information that enables infants to form category representations at the basic and global levels in these studies. The age of the subjects and the nature of the stimuli (i.e., static pictorial instances of the categories) make it improbable that the infants are relying on conceptual knowledge about the "kind of thing" something is to perform successfully in these tasks (cf. Mandler & McDonough, 1993). The studies therefore support the position that both basic and global levels of representation can have a perceptual basis.

Given this state of affairs, at least two important issues remain unresolved. First, is there a sequence to the development of category representations at the two levels in younger infants (i.e., from basic to global or vice versa)? Second, what is the basis of such a sequence? To examine these issues, we have been exploring the emergence of basic-level and global-level category representations in connectionist learning systems. Using as input the dimensions of the stimuli employed in the experiments cited above, we have found that a variety of two layered (i.e., input-output) and three layered (i.e., input-hidden-output) network architectures produce both basic- and global-level category representations and reveal that global-level categories precede basic-level categories in order of appearance. In this paper, we consider the performance of two of these models in detail and examine possible reasons for the observed global-to-basic developmental trajectory.

Simulation 1: Global before Basic

Method

Network Architecture and Training Stimuli. The network had 13 input nodes, 3 hidden nodes and 10 output nodes. Each hidden node received input from all 13 input nodes and in turn sent input to all 10 output nodes. The input nodes encoded 13 attributes of pictorial instances of cats, dogs, elephants, rabbits, chairs, tables, beds and dressers -- stimuli used in the studies described earlier (Behl-Chadha,

in press; Quinn & Eimas, in press-a). These stimuli were realistic color photographs, each displaying an individual mammal or furniture item. They were selected to be nearly the same size as possible so that the infant would use cues other than size as bases for categorization. Three instances of each category were randomly selected as training inputs and an additional instance was randomly selected for generalization testing. Stimulus attributes that served as inputs were head length, head width, eye separation, ear separation, ear length, nose length, nose width, mouth length, number of legs, leg length, vertical extent of the stimulus, horizontal extent of the stimulus, and tail length. The attribute values were measured directly from the stimuli in centimeters and then linearly scaled so that the highest value on any given attribute was 1.0. If a stimulus did not possess an attribute, then the value for that attribute was encoded on its respective input node as 0.0.

This manner of parsing input patterns into component attributes, and using the attribute values along with certain assumptions about processing to make predictions about the formation of category representations, has been used in previous infant categorization investigations (e.g., Strauss, 1979; Younger, 1990). The attributes of horizontal and vertical extent were chosen because of their correspondence with the width and height of the furniture and mammal stimuli. The large number and detailed nature of the facial attributes were selected because of evidence that infants are highly attracted to facial configuration information (e.g., Johnson & Morton, 1991). Young infants also appear to use information from the face and head region of cats and dogs to categorically distinguish between them (Quinn & Eimas, in press-b). For example, infants familiarized with cat stimuli in which only the face and head region was visible (the body information had been occluded), preferred novel dog faces over novel cat faces. However, infants familiarized with cat stimuli in which only the body information was visible (the face and head region was occluded), looked equivalently to novel dog and cat bodies. Subsequent control experiments revealed that the dog preference in the "face and head visible" group could not be attributed to a spontaneous preference for dog faces or to an inability to discriminate among cat faces. Facial information would thus seem to provide infants with a necessary and sufficient basis to form a category representation for cats that excludes dogs. Quinn and Eimas also showed that the cues for this category representation of cats resided in the internal facial region (inclusive of the eyes, nose and mouth) and along the external contour of the head.

Ten output nodes were responsible for indicating the basic and global category identity of the stimuli: cat, dog, elephant, rabbit, chair, table, bed, dresser, mammal and furniture. Each stimulus was associated with two of the ten output nodes, one for the basic level, the other for the global level. Given that the range of activation of the units was from 0.0 to 1.0, the system was considered to have correctly recognized the category identity(ies) (i.e., global and basic) for a given stimulus if it activated the output node(s) associated with that stimulus to a value(s) greater than 0.50 and activated the output nodes corresponding to stimuli from other categories to values less than 0.50.

Training and Testing Procedure. Training consisted of presentation of the 24 stimuli in a random order with replacement for 7200 training sweeps (1 sweep = 1 presentation of a single stimulus). Generalization testing consisted of one presentation of a novel exemplar from each category.

Implementation. The simulation was run on the neural network simulator *tlearn* that makes use of a backpropagation learning algorithm (Plunkett & Elman, in press). The network was trained with a random seed of 47, a learning rate of 0.3 and a momentum of 0.7. Comparable results were obtained with two other random seeds. Thus, while we present data from one random seed in detail, the basic results (with only minor variations) are extendable to a variety of starting seeds (and this is true for each of the simulations reported in the paper).

Results and Discussion

Performance of the network is shown in Figure 1 where the Root Mean Square (RMS) Error (reflecting the discrepancy between actual and correct response for a given input) is plotted as a function of training sweeps. Category learning began at 120 sweeps with the global distinction between mammals and furniture. At 960 sweeps, the elephant exemplars were learned. By 3600 sweeps, the beds, cats, dogs, dressers, rabbits and tables were categorized. Learning was completed at 7200 sweeps when chairs were correctly classified.




Figure 1. Root mean square error of Simulation 1 as a function of training sweeps. Category labels along the sweep axis are positioned to show the categories that have emerged at 120, 960, 3600 and 7200 sweeps.

It is interesting to consider the representations of the input patterns that emerged on the hidden units during training. Figure 2 presents a 3-dimensional plot of the mean activation values on hidden nodes 1, 2 and 3 (relabeled as X, Y and Z) generated by the 8 categories of stimuli at three points in training. Panel A (left) shows that at 8 sweeps, all 8 categories cluster closely together. Panel B (center) shows

that at 480 sweeps, only mammals and furniture were segregated. Finally, Panel C (right) reveals that at 7200 sweeps mammals and furniture were segregated along the z-axis and each basic-level category had its own location

recognized as distinct basic-level categories in this simulation. The results show that the global-to-basic sequence is obtained even without the face and tail information from the mammal stimuli, indicating that it is




Figure 2. Mean activation values of hidden nodes 1, 2 and 3 (relabelled as X, Y and Z) for each category at (a) 8 sweeps, (b) 480 sweeps and (c) 7200 sweeps.

within the "mammal" and "furniture" planes. Figure 2 thus provides an illustrative example of how category structure emerges over time on the representational units.

The results of the simulation are consistent with the findings that young infants can form perceptually based category representations for mammals and furniture at both global and basic levels (Behl-Chadha, in press; Quinn et al., 1993). The complete learning sequence is also consistent with a developmental progression from global- to basic-level category distinctions, a pattern which corresponds with the developmental course of category acquisition in older infants, but with what were presumed to be conceptually based representations for animals and artifacts (Mandler et al., 1991; see also McClelland, McNaughton & O'Reilly, 1995).

One question that arises is whether the early appearance of the global-level representations occurred because the mammals provided activation on the various input nodes devoted to the processing of face and tail information, whereas the furniture stimuli did not. One way of asking this question is: Will the global-to-basic sequence be observed in a model operating on input that does not include the face and tail information? To provide an answer, we repeated the initial simulation, but in this case with just 4 input nodes (number of legs, leg length, vertical extent and horizontal extent).

By 7200 sweeps, the no face-no tail network differentiated the global-level categories; only the basic-level category of dressers had appeared at this point (and if one looks earlier into the training sequence, one finds that dressers were first responded to as furniture and only subsequently as dressers). Basic-level categorization of rabbits and tables (14,400 sweeps), elephants (21,600 sweeps), chairs (28,800 sweeps) and beds (43,200 sweeps) completed the learning sequence. Dogs and cats were not

not simply a consequence of specialized processing for mammals and that it may be generalizable beyond the mammal-furniture distinction. In addition, the inability of the model to categorically differentiate cats and dogs is consistent with the finding that young infants rely on head and face information to distinguish between them (Quinn & Eimas, in press-b).

A second issue raised by the finding of global-to-basic category development is whether the global level would have emerged before the basic level if the network had not been trained at the global level. To answer this question, we repeated the initial simulation (with 13 inputs), but in this case without the two global-level output nodes. There was thus no teaching signal at the global level. While this manipulation prevents us from determining whether the patterns were responded to as mammals and furniture at the output layer, we can still inspect the representation of the patterns at the hidden layer at different points during training.

What is observed in this simulation is that the global level of representation still emerges before the basic level. At 480 sweeps, the mean activation values for cats, dogs, elephants and rabbits on hidden node 1 were 0.191, 0.160, 0.084 and 0.212, whereas those for chairs, tables, beds and dressers were 0.816, 0.795, 0.833 and 0.831. This global-level separation was maintained throughout the remainder of training. In contrast, hidden nodes 2 and 3 at 480 sweeps did not allow for partitioning of inputs into basic-level categories. Consistent with these observations is the timing of emergence of the basic-level categories as assessed by their corresponding output activation values. Elephants were distinguished at 960 sweeps, followed by dogs, rabbits, chairs and dressers at 3600 sweeps, beds at 7200 sweeps and cats and tables at 10,800 sweeps. The results of this simulation are important because they suggest that the early

appearance of global-level categories occurs even when the network is not being trained at the global level. The global level may thus be thought of as a "primary" representation that occurs in the course of mapping a set of structured (but uncategorized) inputs onto basic-level categories.

Simulation 2: Arbitrary Global-level Category Learning

In this section, we examine a possible reason for the global-to-basic developmental sequence. One idea is that global occurs before basic because of the nature of global-level categories. This idea can be tested by orthogonalizing (i.e., crossing) the stimulus dimensions relevant for the global level. That is, one can change the nature of the categories at the global level and determine if the global-to-basic trend still emerges. To this end, we examined the performance of a network taught to assign cats, elephants, chairs and beds to one arbitrary global-level category called A and to respond to dogs, rabbits, tables and dressers as members of a second global-level category called B.

Method

The only change from the first simulation described in the preceding section was that the output node previously coding for mammals was reassigned to code for A stimuli (cats, elephants, chairs and beds) and the output node that earlier coded for furniture now coded for B stimuli (dogs, rabbits, tables and dressers).

and chairs and beds (3600 sweeps) activated the A output node, and concluded with tables recognized as members of the B category (7200 sweeps). Thus, both arbitrary global-level and basic-level categories were learned, but in no particular order.

A more complex picture regarding performance of this simulation emerges when one examines the mean activation values on the 3 hidden nodes for the various categories. Figure 3 presents a 3-dimensional plot of these values at 8, 480 and 7200 sweeps. The 8 sweeps plot (Panel A) reveals no clear partitioning of the categories. However, at 480 sweeps (Panel B), the mammals and furniture have been segregated, indicating that perceptual global-level category structure emerged even when the network was being taught on a different "arbitrary" (i.e., non-perceptual) global-level distinction. At 7200 sweeps (Panel C), the arbitrary global differentiation into categories A and B has appeared. There is also a segregation of mammals and furniture. This figure reveals that the hidden nodes coded for two distinct global levels of representation during the course of training: an initial "perceptually based" global level of mammals and furniture and a subsequent "conceptually based" global level of A and B. Such a finding is consistent with the idea that distinct perceptual and conceptual representations develop for object categories during early development (Mandler & McDonough, 1993).

The major result of this simulation is that manipulating the structure of the global level categories interfered with




Figure 3. Mean activation values of hidden nodes 1, 2 and 3 (reabeled as X, Y and Z) for each category at (a) 8 sweeps, (b) 480 sweeps and (c) 7200 sweeps. In (c) the connected points indicate members of one arbitrary global-level category. Separation between perceptual global-level categories can also be seen.

Results and Discussion

In this simulation, learning occurred at both basic and arbitrary global levels, without one level clearly preceding the other. At the basic level, the order of classification of the training exemplars was: elephants (960 sweeps), rabbits (1920 sweeps), cats and dressers (2760 sweeps), and dogs, tables and beds (3600 sweeps). Chairs, even at 7200 sweeps, failed to elicit a consistent response from the appropriate output node. Learning at the arbitrary global level also began at 960 sweeps with elephants activating the A output node and dogs, rabbits and dressers activating B. Arbitrary global-level classification continued as cats (1920 sweeps)

the global-to-basic order of category development at least for the arbitrary global-level categories. No clear timing difference was observed in the emergence of representations at the basic and arbitrary global levels. However, the perceptual global level (i.e., mammals distinct from furniture) was the first level of category representation to appear, even though the network was not explicitly taught to make this distinction. The overall pattern of performance thus indicates that the nature of global categories (i.e., perceptual vs. arbitrary) is a critical factor in their early appearance during training.

General Discussion

Connectionist accounts of cognitive development are growing in influence (Elman, Bates, Karmiloff-Smith, Johnson, Parisi & Plunkett, in press; McClelland, 1989) and this paper represents an attempt to apply a connectionist analysis to the issue of how category representations arise at basic and global levels during early development. A set of simple network simulations were found to learn categories at basic and global levels and in a global-to-basic sequence.

A striking result of the simulation with arbitrary global-level category training was that both perceptual and arbitrary global-level categories were formed, despite there being no explicit training for the former. This is reminiscent of the views of Mandler & McDonough (1993) who have argued for distinct perceptual and conceptual levels of category representation in human infants. A key difference between our model and the one proposed by Mandler & McDonough is that in our model a single network forms both types of representations, whereas in their view perceptual and conceptual representations are the products of two complementary processes.

Predictions

Models are often judged by the degree to which they can generate interesting experimental predictions. We therefore offer the following:

1. On the assumptions that our results generalize to other categories and that the early-appearing representations that have emerged for global-level categories are stronger than those coding for the later-appearing basic-level categories (cf. Munakata, McClelland, Johnson, & Siegler, submitted), it should be possible to manipulate task parameters that will make categorization more difficult and observe that basic-level representations are affected to a greater degree than those at the global level. For example, reducing the amount of time a stimulus is exposed to an infant decreases its memorability (Fagan, 1974), so reducing the amount of familiarization time for each of a group of exemplars should make categorization more difficult. The reported simulations suggest that basic-level distinctions would be the first to be affected by a moderate decrease in study time per exemplar.

2. Both global- and basic-level category representations have been observed with 3- to 4-month-old infants (e.g., Quinn & Eimas, in press-a). The global-to-basic sequence observed in the simulations would therefore predict that global-level representations should emerge before basic-level ones sometime prior to 3 months of age.

3. It should be possible to train models with one or more deleted input nodes and examine which, if any, category representations fail to emerge. Such manipulations should be helpful in determining which aspects of the input are critical for certain category distinctions. For example, the no face-no tail model predicts that face and tail information is not necessary for making the category distinction between mammals and furniture, a prediction that can be tested on infants with simple alterations to the mammal stimuli. In addition, it should be possible to remove one or more inputs from the face and head region of cats and dogs to predict which are needed to make this basic-level distinction,

predictions than can be tested by presenting altered cat and dog stimuli to infants (cf. Quinn & Eimas, in press-b).

4. The arbitrary global-level category learning simulation indicates that it may be possible to train subjects, either infants or toddlers, to assign stimuli to arbitrary global-level categories. The simulation also predicts that in the context of such training an initial perceptual global level of category representation will precede formation of both basic- and arbitrary global-level categories.

5. It is of interest to examine the representations that emerge in models with larger numbers of hidden nodes. For example, we have observed that when the number of hidden nodes is 8, 9, 10 or 11, the global-to-basic sequence still emerges. In addition, during the course of training, there is a gradual decrease in the proportion of the overall representation that codes for the global level and a gradual increase in the proportion of the overall representation that codes for the basic level. These findings predict that if infants could be repeatedly familiarized with instances of a given category in successive sessions, then there would be a steady transition from global- to basic-level representation. The results also predict that frequently experienced environmental stimuli may tend to elicit basic-level responding (albeit subsequent to global-level responding).

Concluding Comments

In our view, a strength of the approach we have presented is the correspondence between the experimental work on infant categorization and the network simulations. That is, the input to the models were the dimensions of stimuli presented to infants in a series of studies on the development of perceptual categorization in early infancy (Quinn & Eimas, in press-a). Data from the experimental studies were used in decisions about what inputs to present to the models. There is also evidence that infants use correlated attribute information to perform successfully in various kinds of categorization tasks (e.g., Younger, 1990) -- a manner of information processing that is broadly consistent with the way in which neural networks learn information.

A limitation of the input scheme used in these simulations is that the visual representations infants (or adults) use to recognize objects are still unknown. Moreover, we did not use the object parsing schemes advocated in some contemporary models of object recognition (e.g., Biederman, 1987). It therefore becomes important to examine whether implementations of our models with a range of input descriptions would produce comparable results. However, we believe that our basic observations on perceptual category formation will be robust for the reason that all models of object recognition would encode greater similarity between different mammals than between a mammal and an item of furniture, for example. It is this similarity structure in the input that we believe to be important for the results obtained, rather than the details of what elements of the visual array are encoded.

A second possible limitation is that the networks reported in this paper were trained by a backpropagation learning algorithm -- a teaching signal that drives the gradual reduction of error observed in the networks. One can claim that this manner of learning is questionable in the present

context for at least two reasons. First, there are some who maintain that backpropagation is a biologically unrealistic form of learning (e.g., Crick, 1989). Second, there is no external teacher supervising infants in the perceptual kinds of categorization tasks we have attempted to simulate.

We make three observations regarding these points. First, at least one level of category representation, the perceptual global level, was obtained without an explicit teaching signal. Second, backpropagation has not been ruled out as a biologically plausible form of learning. For example, Plunkett (1996) has speculated that backprojecting neurons might be one mechanism by which backpropagation in the nervous system could be accomplished. Third, backpropagation is thought to be one of an equivalence class of learning algorithms with similar computational properties. For example, networks trained with backpropagation in some instances develop the same representations as those produced by more biologically plausible, Hebbian learning algorithms (e.g., Plaut & Shallice, 1993).

In conclusion, we believe that many of the effects we have observed in these simulations would be extendable to other connectionist architectures, input formats, and learning rules (including unsupervised networks). We also believe that the findings of the simulations along with the experimental predictions generated from them represent an important first step toward a research program which combines experimental studies of infant categorization with techniques of connectionist modelling. Such a program may hold promise for developing a formalized account of category formation by young infants.

Acknowledgements

This research was supported by Grant HD 28606 from the National Institute of Child Health and Human Development (PCQ), the UK Medical Research Council (MHJ), the Human Frontier Science Program (MHJ and PCQ) and the British Academy (MHJ and PCQ). We thank Clay Mash for his assistance in preparing Figures 2 and 3. We also thank Peter Eimas, Annette Karmiloff-Smith, Jean Mandler, Denis Mareschal, Frédéric Vallée-Tourangeau, and Steven Young for helpful discussion.

References

- Behl-Chadha, G. (in press). Basic-level and superordinate-like categorical representations in early infancy. *Cognition*.
- Biederman, I. (1987). Recognition by components: A theory of human image understanding. *Psychological Review*, **94**, 115-147.
- Crick, F. H. C. (1989). The real excitement about neural networks. *Nature*, **337**, 129-132.
- Elman, J. L., Bates, E. A., Karmiloff-Smith, A., Johnson, M. H., Parisi, D., & Plunkett, K. (in press). *Rethinking innateness: Connectionism in a developmental framework*. Cambridge, MA: MIT Press.
- Fagan, J. F. (1974). Infant recognition memory: The effects of length of familiarization and type of discrimination task. *Child Development*, **45**, 351-356.
- Johnson, M. H., & Morton, J. (1991). *Biology and cognitive development*. Cambridge, MA: Blackwell.
- Mandler, J. M., Bauer, P., & McDonough, L. (1991). Separating the sheep from the goats: Differentiating global categories. *Cognitive Psychology*, **23**, 263-298.
- Mandler, J. M., & McDonough, L. (1993). Concept formation in infancy. *Cognitive Development*, **8**, 291-318.
- McClelland, J. L. (1989). Parallel distributed processing: Implications for cognition and development. In R. G. M. Morris (Ed.), *Parallel Distributed Processing: Implications for Psychology and Neurobiology* (pp. 8-45). Oxford: Clarendon Press.
- McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, **102**, 419-457.
- Mervis, C. B. (1987). Child-basic object categories and early lexical development. In U. Neisser (Ed.), *Concepts and conceptual development: Ecological and intellectual factors in categorization* (pp. 201-233). Cambridge, UK: Cambridge University Press.
- Munakata, Y., McClelland, J. L., Johnson, M. H., & Siegler, R. S. (submitted). Rethinking infant knowledge: Toward an adaptive process account of successes and failures in object permanence tasks.
- Plaut, D. C., & Shallice, T. (1993). Deep dyslexia: A case study of connectionist neuropsychology. *Cognitive Neuropsychology*, **10**, 377-500.
- Plunkett, K. (1996). Development in a connectionist framework: Rethinking the nature-nurture debate. *The Newsletter of the Center for Research in Language (UCSD)*, **10**, 3-14.
- Plunkett, K., & Elman, J. L. (in press). *Simulating nature and nurture: A handbook of connectionist exercises*. Cambridge, MA: MIT Press.
- Quinn, P. C., & Eimas, P. D. (in press-a). Perceptual organization and categorization in young infants. In C. Rovee-Collier & L. P. Lipsitt (Eds.), *Advances in infancy research, Vol. 10*. Norwood, NJ: Ablex.
- Quinn, P. C., & Eimas, P. D. (in press-b). Perceptual cues that permit categorical differentiation of animal species by infants. *Journal of Experimental Child Psychology*.
- Quinn, P. C., Eimas, P. D., & Rosenkrantz, S. L. (1993). Evidence for representations of perceptually similar natural categories by 3-month-old and 4-month-old infants. *Perception*, **22**, 463-475.
- Strauss, M. S. (1979). Abstraction of prototypical information by adults and 10-month-old infants. *Journal of Experimental Psychology: Human Learning and Memory*, **5**, 618-632.
- Younger, B. A. (1990). Infants' detection of correlations among feature categories. *Child Development*, **61**, 614-620.