JOURNAL OF EXPERIMENTAL CHILD PSYCHOLOGY **66,** 236–263 (1997)
ARTICLE NO. CH972385

The Emergence of Perceptual Category Representations in Young Infants: A Connectionist Analysis

Paul C. Quinn

Washington and Jefferson College

and

Mark H. Johnson

MRC Cognitive Development Unit

There has been recent interest in the idea that principles governing learning in connectionist networks can form the basis for an alternative understanding of developmental processes (Elman, Bates, Karmiloff-Smith, Johnson, Parisi, & Plunkett, 1996). The present paper can be viewed as a case example of the usefulness (and limitations) of connectionist modeling for the study of infant cognition. Specifically, the paper reports on a series of connectionist models designed to analyze the factors responsible for the emergence of global-level and basic-level category representations in young infants. The models (1) simulated the formation of global-level and basic-level representations, (2) revealed a global-to-basic order of category emergence, (3) uncovered the formation of two distinct global-level representations—an initial ''self-organizing'' perceptual global level and a subsequently ''trained'' arbitrary (i.e., nonperceptual) global level, and (4) displayed a gradual transition from perceptual global-level to perceptual basic-level representation with increasing exposure to training stimuli. Hypotheses for empirical investigations of category development in infants that follow from the modeling efforts are discussed.

A number of investigators interested in early cognitive development have been examining the origins and development of complex category representa-

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 British Academy (MHJ and PCQ), and the Human Frontiers Scientific Program (MHJ and PCQ). P.C.Q. thanks M.H.J. and the Medical Research Council Cognitive Development Unit, London, UK, for hosting him during a sabbatical year—the period when this work was conducted. Both authors thank Clay Mash for assistance in preparing Figs. 4 and 5, and Peter Eimas, Annette Karmiloff-Smith, Denis Mareschal, Jean Mandler, Gary Marcus, Fred Valee-Tourangeau, and two anonymous referees for valuable comments on an earlier version of the paper. We also thank Daphne Maurer, Andrew Oliver, Kim Plunkett, Edmund Rolls, Michael Tarr, and Steven Young for helpful discussion. Requests for reprints should be sent to Paul C. Quinn, Department of Psychology, Washington and Jefferson College, Washington, PA 15301, e-mail: pquinn@vms.cis.pitt.edu.

0022-0965/97 $25.00

tions during the first two years of life (e.g., Mandler, Bauer, & McDonough, 1991; Mervis, 1987; Quinn, Eimas, & Rosenkrantz, 1993). Empirical efforts have been focused on 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). There has also been concern with whether early basic-level representations cohere to form global-level representations or whether basic-level representations evolve from original global-level representations. Much of this work has been in response to the theory of Rosch and Mervis which suggested that categories were initially formed at the basic level and that superordinate categories developed later when the infant grouped together separate basic-level representations (Rosch, Mervis, Gray, Johnson, & Boyes-Braem, 1976; but see Keil, 1979, and Mandler & Bauer, 1988, for evidence to the contrary; see also Mervis, 1987, for a revision of the original theory).

One series of studies relevant to these issues has shown that young infants participating in the familiarization–novelty preference procedure can form category representations at both basic and global levels of exclusiveness (reviewed in Quinn, in press; Quinn & Eimas, 1996b). At the basic level, for example, 3- to 4-month-olds familiarized with domestic cats will generalize their familiarization to novel cats, but dishabituate to birds, horses, dogs, tigers, and even female lions (Eimas & Quinn, 1994; Eimas, Quinn, & Cowan, 1994; Quinn & Eimas, 1996a; Quinn, Eimas, & Rosenkrantz, 1993). The data provide evidence that the infants can form a category representation for cats that includes novel cats, but excludes exemplars chosen from a variety of related basic-level categories. Behl-Chadha (1996) has extended these findings to human-made artifacts by showing that 3- to 4-month-olds can also form individuated representations for chairs and couches each of which exclude 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 their familiarization to novel mammal categories (e.g., deer), but dishabituated to instances of birds, fish, and furniture (Behl-Chadha, Eimas, & Quinn, 1995; Behl-Chadha, 1996). These results indicate that the infants can form a global-level representation of mammals that includes novel mammal categories, but excludes instances of nonmammalian 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 global-level representation for furniture that includes beds, chairs, couches, cabinets, dressers, and tables, but excludes the mammals mentioned above (although possibly not vehicles). 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 questions remain at issue. First, what representations might be utilized by young infants in the formation of categories? Second, on the basis of the data now available, can we predict the course of category development in even younger infants? That is, do the original category representations formed by infants have relatively broad, global extensions or are they more narrowly tuned—perhaps nearly adult ''basic'' in their range of exclusiveness?

To examine these issues in a more formal way than has been done in the past, we have been exploring the emergence of basic-level and global-level category representations in connectionist learning systems. Using as input the measured dimensions of stimuli employed in the familiarization–novelty preference experiments cited above and an input scheme that corresponds with one that has been gaining acceptance among investigators of object recognition (Zhu & Yuille, 1996), we report that a series of connectionist networks relying on a three-layered network architecture (i.e., input → hidden layer → output) produce both basic-level and global-level category representations, and that global-level categories (e.g., mammals, furniture) usually precede basic-level categories (e.g., cats, tables) in order of appearance. After presenting a brief review of the major features of the connectionist approach, we consider the performance of these models in detail and examine the reasons for global-level and basic-level category formation and the observed global-to-basic learning sequence.

Since the publication of McClelland and Rumelhart's *Parallel Distributed Processing, Explorations in the Microstructure of Cognition, Volumes 1 and 2* (McClelland & Rumelhart, 1986a, 1986b), connectionist models have been gaining in influence and are now viewed by many as a promising level of analysis by which to explain cognition and its development (e.g., Clark, 1993; Karmiloff-Smith, 1992a, 1992b; McClelland, 1989; Plunkett & Sinha, 1992). While connectionist models are not incompatible with symbolic models of cognitive functioning and some have argued that complementarities may lead to the development of hybrid models (Clark & Karmiloff-Smith, 1993; Mandler, in press), others have hinted that connectionist models may eventually replace symbolic accounts of cognition (see, for example, the discussion in Smolensky, 1988). Connectionist models may in some instances bring forth levels of detail and precision not present in more classical, introspectionist, verbal descriptions of behavior. Such details may be critical to the eventual realization of a theory of the microstructure of cognition.

Connectionist models have as their basic building blocks neuron-like entities called processing units that compute by way of connections with each

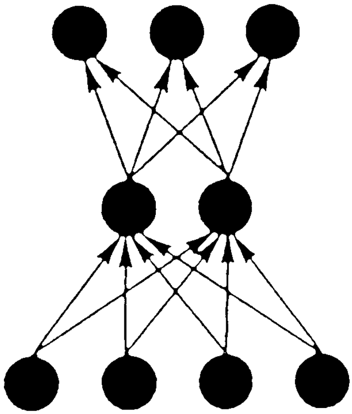

FIG. 1. An example of a 3-layered connectionist architecture showing input, hidden, and output units and the pattern of connectivity between them.

other (see, Crick, 1989, however, for discussion of how processing units are unlike neurons). In most connectionist models, the processing units are hierarchically organized into several layers including an input layer, one or more hidden layers for purposes of internal representation, and an output layer. The architecture of the model is in part determined by the pattern of connectivity of the units within and between the levels. An example of a fully connected input → hidden layer → output network architecture is shown in Fig. 1.

Representations in connectionist models are patterns of activation over the units in the network. As shown in Fig. 2, a single unit's activation is based on two components: the net input to the unit and the activation function of the unit. The net input to a unit is calculated as a weighted sum of the inputs to the unit from the environment or other units. The contribution of each unit to the net input is weighted by the strength of the connection from contributing to receiving unit. The activation function is then used to determine the resulting activity of the unit given the net input—activity that will then be passed forward along output connections.

Processing in a connectionist model occurs through changes in the units' patterns of activation over time. These changes are in turn dependent on changes in the strengths of the connections between units. Connection strengths in many neural networks are initially random values and change as a function of experience with structured input according to a learning algorithm. We used a learning rule known as backpropagation or the generalized delta rule (Rumelhart, Hinton, & Williams, 1986). Backpropagation can be used when the job of the network is to map a set of inputs onto a set of outputs (as is the case in categorization tasks) and was devised to work in

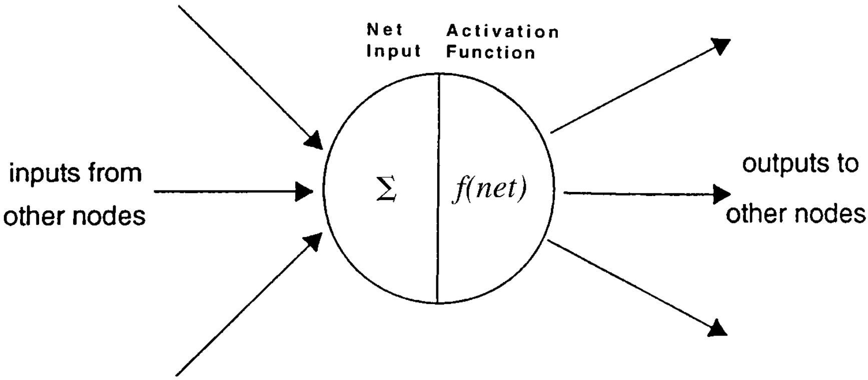

FIG. 2. An illustration of the flow of activation into and out of a unit and how the net activation of the unit is dependent on (1) the net input to the unit and (2) the activation function, *f*(net), of the unit.

particular with networks that have at least one layer of hidden units. The task facing the network is to move from a starting point of arbitrary random weights to an eventual configuration of weights that produces the desired output activations. During the course of training, each time an input is processed through the model to produce a pattern of output activations, these activations are compared with the desired output activations. When differences between actual and desired activations occur on any of the output units, the connection strengths coming into those units are gradually adjusted in directions that reduce the error. The error signals for the various output units are then backpropagated through lower layers of hidden units so that error reducing adjustments in connections strengths can be made throughout the network.

From a developmental standpoint, connectionist models have appeal because they are composed of a small set of simple processing mechanisms from which both qualitative and quantitative predictions can be generated (Elman, Bates, Johnson, Karmiloff-Smith, Parisi, & Plunkett, 1996). Connectionist models also have the advantage that they can account for apparent discontinuities in the development of cognitive abilities without recourse to qualitative changes in processes or representations (McClelland, 1989; Plunkett, Sinha, Moller, & Strandsby, 1992; but see Raijmakers, van Koten, & Molenaar, 1996). Recent connectionist modeling efforts have begun to suggest that mechanisms that differ in kind across development are not necessary to explain, for example, the development of object permanence (Mareschal, Plunkett, & Harris, 1995; Munakata, McClelland, Johnson, & Siegler, in press) and as a result have begun to change the conventional wisdom about the course of human cognition (Elman et al., 1996). However, whether it will be the case that connectionist models based on principles of developmental continuity will adequately describe all facets of concept formation remains an open question (Karmiloff-Smith, 1992a, 1992b; Mandler, in press). The present paper can be viewed as part of an effort to explore the usefulness of connectionist modeling in explaining an emerging and important developmental finding—the formation of basic-level and global-level category representations by young infants. We point out that while connectionist models have a history of being used to model categorization (e.g., Feldman & Ballard, 1982; Knapp & Anderson, 1984), little work has focused on the development of categorization (for exceptions, see Miikkulainen & Dyer, 1991, Rumelhart & Todd, 1993, and Schyns, 1991), and none has attempted to model the infant's perceptual categorization abilities.

Our primary aim in this paper is to illustrate the benefits and limitations of applying connectionist modeling to a well-studied question in early perceptual–cognitive development. We do this by presenting a number of simple connectionist models that allow us to explore the influence of intrinsic architecture and extrinsic input structure on the formation of category representations. We then discuss the extent to which the findings of our simulations correspond to experimental data on categorization obtained from infants. In taking this approach we do not assume that the category representations formed by humans are merely a reflection of environmental structure, or the product of a single learning mechanism. Rather, the category representations that emerge in the simulations to follow are viewed by us as outcomes of the interaction between a particular input structure and network architecture. The category training sequences that appear in the networks should thus not be construed as the actual time course of category learning adhered to by infants, but as examples of what a certain class of connectionist models predict about the developmental course of category emergence given a specific input structure. In our view, the major contribution of the simulations lies in the hypotheses they generate, and relation to empirical data from infants.

SIMULATIONS PART I: GLOBAL BEFORE BASIC

Method

Network architecture and training/test stimuli. To begin, a model was developed to examine learning of basic-level and global-level category representations for instances of furniture and mammals. The model was a network with 13 input nodes, 3 hidden nodes, and 10 output nodes. The input nodes encoded 13 attributes of pictorial instances of cats, dogs, elephants, rabbits, beds, chairs, dressers, and tables. The cat and dog stimuli were used in Eimas et al. (1994), Quinn et al. (1993), and Quinn and Eimas (1996a), and the elephant, rabbit, and furniture stimuli were used in the studies of Behl-Chadha (1996) and Behl-Chadha et al. (1995). 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 (i.e., cues believed to be more nearly category defining) as possible bases for categorization. Three instances of each category were randomly selected to be inputs during training. An additional instance of each category was randomly selected for the test of generalization.

Attributes of the stimuli that served as inputs were as follows: head length, head width, eye separation, ear separation, ear length, nose length, nose width, mouth length, number of legs, leg length, vertical extent of mammal bodies and furniture stimuli (exclusive of leg length), horizontal extent, and tail length. The attributes were measured directly from the stimuli in centimeters and then linearly scaled so that the highest value on each attribute was 1.0. The scaling procedure involved normalizing the values of a given attribute by dividing each value by the largest value of that attribute. Scaled values were used instead of actual values of the attributes because we were concerned that performance of the model might be unduly influenced by attributes with the largest input values. If a stimulus did not possess a particular attribute, then the value for that attribute was encoded on its respective input node as 0.0. Actual activation values assigned to the stimulus attributes for each training and test pattern will be supplied by the corresponding author upon request.

Parsing the 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 investigations of infant categorization (e.g., Sherman, 1985; Strauss, 1979; Younger, 1990). The input attributes can be divided into two classes—those that encoded geometric aspects of the stimuli and those that encoded aspects of the face region of the mammal stimuli. The coding of geometric aspects of the stimuli corresponds with the ''skeleton extraction'' model of object recognition recently proposed by Zhu and Yuille (1996). In referring to geometric input attributes with labels such as ''number of legs'' and ''leg length,'' we do not mean to imply that the infant has a conceptual understanding of such attributes. The input attributes were all measurable dimensions of the surface properties of the stimuli, and as such, potentially available to low-level visual parsing routines that segment a skeleton outline of a shape's silhouette into a number of component attributes.

The large number and detailed nature of attributes from the facial region were selected on the basis of evidence that infants are highly attracted to facial configuration information (e.g., Johnson & Morton, 1991). There are also psychophysical and neurophysiological data suggesting that at least some of the face and head attributes in the input scheme may be used in face recognition (Rhodes, 1988; Yamane, Kaji, & Kawano, 1988; Young & Yamane, 1992). Furthermore, there are data indicating that young infants use information from the face and head region of cats and dogs to categorically distinguish between them (Quinn & Eimas, 1996a). 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 the 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-level and global-level category identity of the stimuli: cat, dog, elephant, rabbit, bed, chair, dresser, table, mammal, and furniture. Each stimulus was associated with 2 of the 10 output nodes, one for the basic level, the other for the global level. Given that the range of activation of the units in the network was from 0.0 to 1.0, the network was considered to have correctly recognized the basic-level or global-level category identity for a given stimulus if it activated the basic-level or global-level output node associated with that stimulus to a value greater than 0.50 and activated the output nodes corresponding to stimuli from other categories to values less than 0.50.

Three hidden nodes were chosen on the basis that this would be the minimum number needed to represent 8 different categories at the basic level. Each hidden node received input from all 13 input nodes and each hidden node in turn sent output to all 10 output nodes. The purpose of the hidden nodes is to re-represent (and in this case to compress) the information from the input patterns into an efficient coding scheme.

Training and testing procedure. Training consisted of presentation of the 24 stimuli in a random order with replacement (as determined by a random seed) for 7200 training sweeps (one sweep equal to one presentation of a single stimulus pattern). Each stimulus was presented to the network by feeding in its attribute values to the appropriate units across the input layer. Testing for generalization to novel members of the training categories consisted of one presentation of a novel exemplar from each category.

Implementation. The simulations were run on the neural network simulator called tlearn (Plunkett & Elman, 1997). tlearn makes use of the back-propagation learning algorithm described earlier (for a more detailed explanation, see Rumelhart et al., 1986). The network simulation reported in this section of the paper was trained with a random seed (*RS*) of 47, a learning rate (*LR*) of 0.3, and a momentum (*M*) of 0.7. The values of *LR* and *M* were chosen through pilot simulations because they yielded optimal category learning. It is important to note that each simulation reported was conducted with two additional random seeds and in each case the same

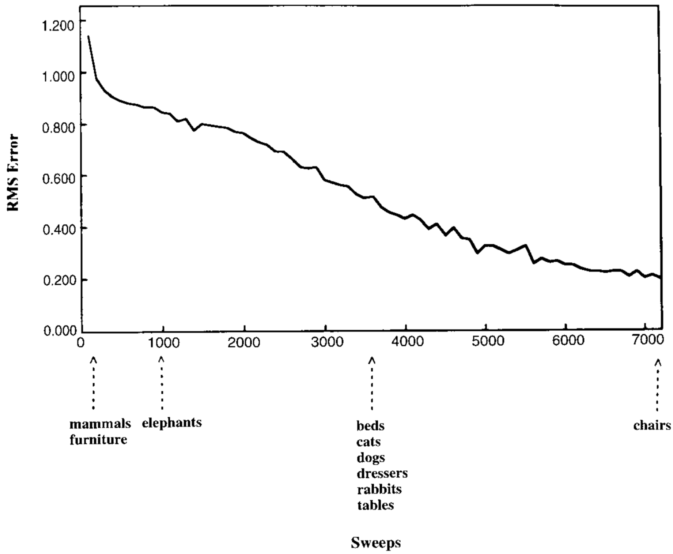

FIG. 3. Root mean square (RMS) error of the initial network reported in Part I as a function of training sweeps. Category labels along the sweep axis are positioned to show the categories that have emerged at 120, 960, 3600, and 7200 sweeps.

overall results were obtained with only minor variations. Thus, while we present data from one random seed in detail, the basic results are extendable to a variety of starting seeds.

Results and Discussion

Performance of the network is shown in Fig. 3. In this figure, a measure of global error known as the root mean square error or RMS (which reflects the discrepancy between actual and correct response to a given input) is plotted as a function of training sweeps. As can be seen, performance of the model improved with increases in the number of training sweeps. Error reduction proceeded quickly over the first few hundred sweeps and was more gradual thereafter.

Category learning was initially manifest at 120 sweeps with the global-level distinction between mammals and furniture emerging for both training and test stimuli. At 960 sweeps the elephant training exemplars were learned, although the novel elephant did not yet elicit generalization. By 3600 sweeps, the cat, dog, rabbit, and elephant training and test exemplars were all categorized. The network also performed correctly on table, bed, and dresser training exemplars, and generalized appropriately to novel instances of each. Learning was completed at 7200 sweeps when the training and test instances of chairs were correctly categorized. The results of the simulation are consistent with the findings that young infants, 3 to 4 months of age, can form perceptually based category representations for mammals and furniture at both global and basic levels (Behl-Chadha, 1996; Quinn & Eimas, 1996b). The complete learning sequence is also consistent with a developmental progression from global category distinctions to more basic ones, a pattern of learning 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 (e.g., Mandler et al., 1991; Mandler & McDonough, 1993).

It is interesting to consider the representations of the input patterns that emerged on the hidden units. The activity of the hidden units can best be summarized by examining the mean activation values corresponding to each category at different points during training. Figure 4 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 two points near the beginning of training and a third point at the end of training. Across the plots, each category corresponds to a point which moves within a stationary three-dimensional cube during the course of network learning as the internal representation of the category changes. 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 within the ''mammal'' and ''furniture'' planes. Figure 4 thus provides an instructive example of how category structure emerges over time on the representational units.

Another connectionist model that has revealed a similar time course of category learning from general to more specific levels is that of Rumelhart and Todd (1993; also discussed in McClelland, McNaughton, & O'Reilly, 1995). Plants and animals were learned before trees and birds which were learned before oaks and robins. However, this model discovered ''conceptual'' structure by learning a set of propositional statements about concepts, ''rather than by percepts that directly provide some information about the concepts'' (McClelland et al., p. 428). For example, the network learned that a canary ''is living,'' ''can grow,'' ''has skin,'' and ''can sing,'' whereas an oak tree ''is living,'' ''can grow,'' ''has bark,'' and ''has roots.'' These attributes are conceptually rich, and we cannot assume that a young infant would be able to use them in initial encounters with members of the animal and plant categories. Thus, a critical contribution of our simulation is to show that the general to specific trend in the development of category representations can be obtained even when the network is operating on the kind of perceptually based input that is presumably available to a young infant.

An additional connectionist simulation that has produced a general to specific trend in the development of category representations is that of Schyns

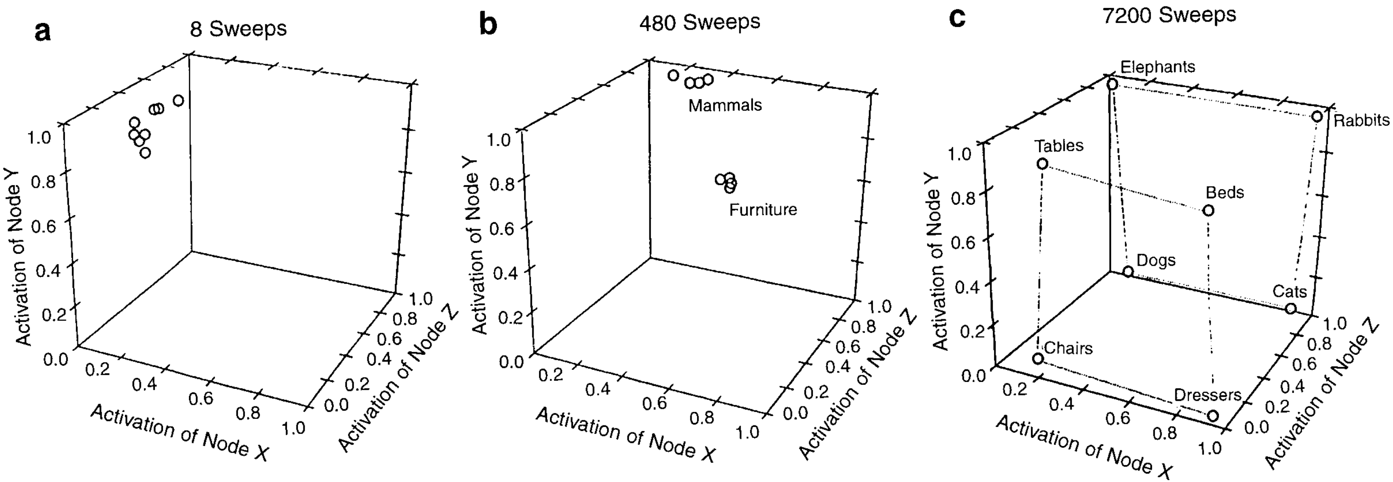

FIG. 4. Mean activation values of hidden nodes 1, 2, and 3 (labeled X, Y, and Z, respectively) for each category at (a) 8 sweeps, (b) 480 sweeps, and (c) 7200 sweeps. At 8 sweeps, the mean activation values for nodes X, Y, and Z for all eight categories were 0.13, 0.88, and 0.40. At 480 sweeps, the mean activation values for nodes X, Y, and Z for mammals were 0.22, 0.95, and 0.97; for furniture the values were 0.81, 0.92, and 0.04.

(1991). While Schyns did not directly address the question of the developmental ordering of global-level vs basic-level categories, he did use a Kohonen (1984) self-organizing network to show that basic-level categories such as ‘‘bird’’ were learned before subordinate-level categories such as ‘‘robin’’ and ‘‘crow.’’ However, the stimuli that Schyns used as input to represent ‘‘natural’’ categories such as ‘‘bird’’ and ‘‘dog’’ were in fact categories constructed from dot patterns that were distortions of arbitrary prototypes (Posner & Keele, 1968). While this work helps us understand the similarity relations discovered by an unsupervised learning algorithm, it is less informative about how the categorization routines of infants might perform on realistic exemplars of naturally contrasting categories.

SIMULATIONS PART II: NO FACE–NO TAIL NETWORK

In the initial simulation the coarse category distinction between mammals and furniture was learned before the more fined-grained distinctions at the basic level. This may simply be because the model responded to the fact that the mammals provided activation on the various input nodes devoted to the processing of face and tail information, whereas the furniture stimuli did not. That is, activation of the face and tail inputs (or lack thereof) provided the basis for the mammal-furniture distinction (see, Rakison, 1996, and Rakison & Butterworth, in press, for evidence that infant responding to global category distinctions may be based on the representation of salient part differences). For this reason, it is of interest to learn how the network will perform without information from the face and tail regions of the mammal stimuli. We therefore conducted an additional simulation with a network containing 4 input units, 3 hidden units, and 10 output units. For this model, training and testing occurred with the same stimuli used in the initial simulation, but in this case the inputs were restricted to number of legs, leg length, vertical extent, and horizontal extent, thereby insuring that no zero values occurred among the input attributes. The network was trained with the random seed used in the first simulation, a similar learning rate of 0.2 and a smaller momentum of 0.3. Different parameter settings were used in the present simulations to accommodate the smaller number of input nodes.

Category learning began in the ‘‘No Face–No Tail’’ simulation at 3600 sweeps with dressers responded to as furniture, but not as dressers. By 7200 sweeps, the global-level categories (both training and test instances) were differentiated; only the basic-level category of dressers had appeared by this point. Basic-level category recognition of training and test instances 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 recognized as distinct basic-level categories in this simulation. This model thus learned the global-level category identity of the entire set of input and test patterns, but learned the basic-level category identity for only a subset of the patterns—a result indicating that the global-level advantage is observed even when face and tail information from the mammal stimuli is not provided as input. The global-level advantage and the failure to distinguish dogs and cats were also obtained with two additional random seeds. The outcome of these simulations corresponds well with the findings of Quinn and Eimas (1996a) who showed that young infants require information from the head and face region to make the basic-level category distinction between cats and dogs. The global-level category precedence observed in these simulations is moreover suggestive that the global-to-basic sequence is not *entirely* a consequence of specialized processing for mammals and that it may be generalizable beyond the mammal–furniture distinction.

When comparing the simulations reported in Parts I and II (those performed with and without face and tail information), one may note that the acquisition of the global-level categories occurred by 120 and 7200 sweeps with and without face and tail information, respectively. That is, acquisition of the global-level categories took 60 times longer when the face and tail features were not present. One may be tempted to conclude from this result that categorization of mammals and furniture by networks (and possibly by infants) is more difficult without face and tail information. While we favor this interpretation (see discussion in Simulations Part III below), one needs to keep in mind that the two simulations are not directly comparable because the simulation conducted without face and tail information was performed with lower values for learning rate and momentum (to accommodate the smaller number of input nodes). As such, at least some of the rate difference in the learning of global-level categories in the two networks may be attributable to the smaller parameter settings. The most reasonable conclusions from Simulations Part I and II would thus appear to be that global-level category differentiation (1) is still possible without face and tail information, but (2) may be more difficult and thus proceed more slowly without face and tail information.

SIMULATIONS PART III: NETWORK WITHOUT GLOBAL-LEVEL CATEGORY TRAINING

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 to assign each of the various inputs to either the mammal or furniture global-level categories. To answer this question, we repeated the initial simulation reported in Part I, 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. It should be noted that in the simulations with global-level category output nodes and hidden nodes, the emergence of global-level categories as measured by activation values on the global-level category output nodes occurred when there was differentiation of these categories on at least one of the hidden nodes. We therefore believe that inspection of hidden node activations is a reasonable way of assessing representation of global-level categories in a network without global-level category output nodes.

What is observed in the model 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. Basic-level categories were learned later in the training sequence (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 (replicated with two other random seeds) 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 might thus be thought of as a ''primary'' representation that occurs in the course of mapping a set of categorically structured inputs onto eventual basic-level representations.

A question that arises based on the simulations already reported in Parts II and III is whether the global level would emerge before the basic level without global-level category training and without face and tail information. To answer this question, another simulation was performed with a network containing only 4 input units (number of legs, leg length, vertical extent, and horizontal extent), 3 hidden units, and 8 output units (one corresponding to each of the basic-level categories). This network was trained with the same random seed and parameter values used in the simulation reported in Part II. Learning in the network began with beds and dressers (7200 sweeps), continued with tables (14,400 sweeps), rabbits and chairs (21,600 sweeps), and concluded with elephants (28,800). Dogs and cats were not differentiated by this network. Perhaps more importantly, global-level category differentiation of mammals and furniture did not emerge on any of the hidden units during the course of training. What this result indicates is that the early appearance of the mammal and furniture global-level categories in the initial simulation conducted in Part III, that performed without global-level training, is likely due to presence vs. absence of face and tail attribute information. Such network performance is consistent with recent empirical work indicating that infants can form global-level category representations, but only when the exemplars are presented so as to preserve salient attribute differences between the categories (Rakison, 1996; Rakison & Butterworth, in press).

Thus far in the paper, we have presented (1) a series of simulations with different architectures and training stimuli that resulted in learning of global-level categories before basic-level ones (with one exception), an outcome of theoretical significance given the traditional basic-to-superordinate view (Rosch et al., 1976), and (2) gained some insight into reasons for the early appearance of the global level (e.g., salient attribute differences between global-level categories). In the next two sections, we further explore possible reasons for the global-to-basic learning sequence.

SIMULATIONS PART IV: ARBITRARY GLOBAL-LEVEL CATEGORY LEARNING

An idea hinted at in Parts II and III is that global occurs before basic because global-level category formation requires a coarser ''cut'' of the input dimensions than does basic-level category formation. By this view, global-level categories are learned as an initial step on the path to basic-level category learning. 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 from perceptual to arbitrary and determine if the global-to-basic trend still emerges. To this end, we examined the performance of two networks taught to assign cats, elephants, chairs, and beds to one arbitrary global-level category which we will call A and to respond to dogs, rabbits, tables, and dressers as members of a second arbitrary global-level category called B.

Method

The network architecture, stimuli, and training/testing procedures (including parameters) were the same as those used in the initial round of simulations reported in Part I. The only change 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). The major results, namely, the difficulty of learning arbitrary global-level categories and the early appearance of perceptual global-level categories, were observed with 2 additional random seeds.

Results and Discussion

In the model, learning occurred at both basic and arbitrary global levels, but without one level clearly preceding the other. At the basic level, the order of classification of the training exemplars was as follows: elephants (960 sweeps), rabbits (1920 sweeps), cats and dressers (2760 sweeps), and dogs, tables, and beds (3600 sweeps). The training instances of chairs, even at 7200 sweeps, failed to elicit a consistent response from the appropriate output node. Learning at the arbitrary global level was also first observed 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) 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, in this network, both arbitrary global-level and basic-level categories were learned, but in no particular order.

A more complex picture regarding performance of the model emerges when one examines the mean activation values on the 3 hidden nodes for the various categories. Figure 5 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 8 categories. However, at 480 sweeps (Panel B), the mammals and furniture have been segregated. This result indicates that perceptual global-level categories emerged even when the network was being taught on an arbitrary (i.e., nonperceptual) global-level distinction—a finding consistent with the perceptual global-level emergence in the earlier model conducted with the full set of inputs, but without global-level output nodes. The mean activation values for each category at 7200 sweeps (Panel C) have been connected in two ways to show that while the perceptual global-level category distinction between mammals and furniture was preserved (top display), the arbitrary global-level differentiation into categories A and B also emerged (bottom display). This analysis indicates that the hidden nodes have coded for two distinct global levels of representation using exactly the same mechanism for both: an initial perceptually based global level and a subsequent arbitrary global level. The extent to which the latter might relate to a conceptually based global level (cf. Mandler, 1997) is examined in the General Discussion.

The major finding from this simulation is that changing the nature of the global level categories interfered with the global-to-basic order of category development at least for the arbitrary global categories. No clear timing difference was observed in the emergence of representations at the basic level and arbitrary global level. However, it was of interest to again find that 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. It should be acknowledged that the early appearance of the perceptual global level in this simulation was dependent on providing the network with the entire set of input attributes. When the global crossover simulation was repeated using the no face–no tail input scheme, the perceptual global level did not emerge. The overall pattern of results indicates that the nature of global categories (i.e., perceptual vs arbitrary) is a critical factor in their early appearance in the simulations.

The fact that perceptually based global-level representations for mammals and furniture emerge first in a model in which the perceptually-based global level was not explicitly taught lends further support to the suggestion that perceptual global-level categories may be a necessary intermediate representation on route from structured (but uncategorized) inputs to perceptual basic-level categories (under certain input schemes). This suggestion could be fur-

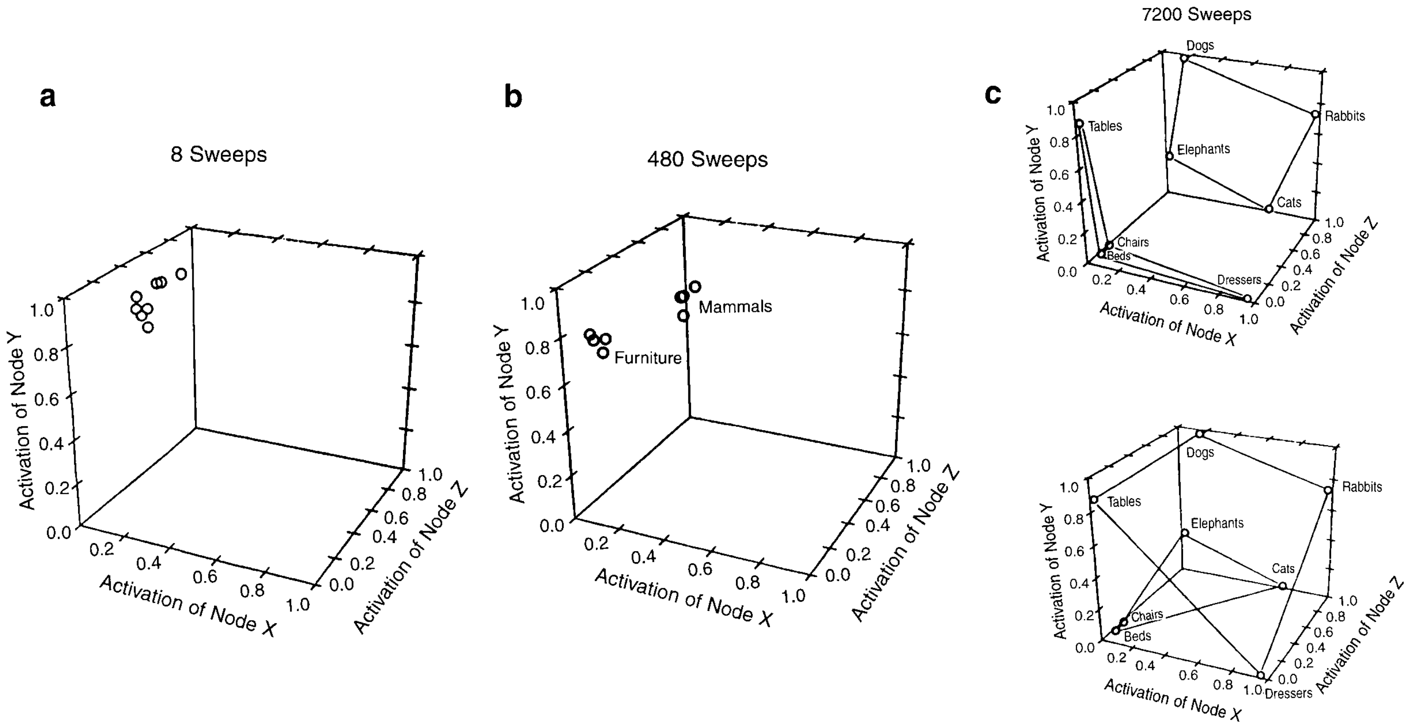

FIG. 5. Mean activation values of hidden nodes 1, 2, and 3 (labeled X, Y, and Z, respectively) for each category at (a) 8 sweeps, (b) 480 sweeps, and (c) 7200 sweeps. At 8 sweeps, the mean activation values for nodes X, Y, and Z for all eight categories were 0.13, 0.88, and 0.40. At 480 sweeps, the mean activation values for nodes X, Y, and Z for mammals were 0.07, 0.67, and 0.86; for furniture the values were 0.10, 0.77, and 0.09. At 7200 sweeps, connections between categories have been drawn to emphasize the natural global distinction between mammals and furniture (top display) and the nonnatural global distinction between categories A and B (bottom display).

ther strengthened if it could be shown that the perceptual global level does not appear on route to the arbitrary global level when there is no training on basic-level categories—a learning sequence that can be tested by running the Global Crossover simulation again, but in this instance without the basic-level category output nodes. The model thus consists of 13 input nodes, 3 hidden nodes, and 2 output nodes, one for category A, the other for category B. The simulation was run with the same training parameters and random seeds as those used in the Global Crossover case. In each instance, the arbitrary global-level categories A and B were learned by 7200 sweeps. Notably, at no time during training did any of the 3 hidden nodes code for the perceptual global-level distinction between mammals and furniture. The results are thus consistent with the idea that with naturally clustered input sets, perceptual global-level categories emerge early and automatically, and as necessary precursors to perceptual basic-level categories.

It also does not appear to be the case that perceptual global-level categories are necessary precursors of arbitrary basic-level categories. Although natural basic-level categories have been presumed to be perceptually based (e.g., Rosch et al., 1976), others have sought to demonstrate the existence of ''basic-like'' category representations for objects that have no clear perceptual basis (e.g., the ad hoc categories of Barsalou, 1983). We therefore examined whether a network would form perceptual global-level categories on route to arbitrary basic-level categories (i.e., categories that violate the correlational structure of the environment). The network had 13 input nodes, 4 hidden nodes, and 6 arbitrary basic-level output nodes. Each arbitrary basic-level category consisted of 2 mammals and 2 furniture items (Category 1: Cat1, Rabbit3, Chair2, Bed1; Category 2: Dog3, Elephant2, Table1, Dresser1; Category 3: Cat2, Elephant1, Chair1, Table3; etc.). After 14,400 training sweeps, the model learned all 6 arbitrary basic-level categories, but at no time did any of the hidden nodes code for the perceptual global-level category distinction between mammals and furniture. It thus does not appear that perceptual global-level categories are formed on route to arbitrary basic-level categories. The evidence continues to point to the idea that perceptual global-level categories may be formed only on the path to perceptual basic-level categories.

SIMULATIONS PART V: NUMBER OF HIDDEN NODES

As a second line of inquiry into the global-to-basic developmental trend, we examined the relation between the number of hidden nodes and the order in which category representations emerge in the network. In general, hidden nodes represent combinations of input attributes and potentially more abstract aspects of the input patterns. In this sense, hidden nodes are said to create internal representations of the input patterns. The presence of one or more hidden nodes is in fact critical to solving certain types of problems in which the similarity structure of the input differs greatly from that of the output (Minsky & Papert, 1969; Rumelhart et al., 1986). For example, hidden nodes are necessary to solve the classic XOR logic problem because they can represent conjunctions of inputs and thereby change the similarity structure of those inputs sufficiently to allow a solution to be learned.

Thus far in the paper, the global-to-basic results we have reported emerged from a network architecture with 3 hidden nodes. It is possible, however, that at some higher number of hidden nodes, the order in which global-level and basic-level representations emerge may be different. In the final set of simulations, we investigated whether the global-to-basic learning sequence would hold up when the number of hidden nodes matched or exceeded the total number of categories to be represented.

Method

The only change in the network architecture compared with that used in the initial round of simulations reported in Part I was to increase the number of hidden nodes to 10 and 11 (given that 2 global + 8 basic = 10 total categories). The stimuli, training parameters, and generalization testing procedures remained unchanged, and the major result (i.e., global-to-basic categorization) was obtained with 2 additional random seeds.

Results and Discussion

In the network with 10 hidden nodes, classification began with the distinction between mammals and furniture at 120 sweeps, followed by cats, elephants, and tables (480 sweeps), rabbits, dogs, and dressers (960 sweeps), and beds and chairs (3600 sweeps). Performance of the network with 11 hidden nodes was comparable with differentiation of mammals and furniture again emerging at 120 sweeps and basic level categories appearing thereafter. Both networks provide evidence that the global-to-basic sequence of category learning does not depend on a specific number of hidden nodes.

We further questioned whether the global-to-basic learning sequence would hold up in a network in which the number of hidden nodes matched the number of basic-level category output nodes and there were no global-level category output nodes. We therefore trained a network with 13 input nodes, 8 hidden nodes, and 8 output nodes (one for each of the 8 basic-level categories). The simulation was conducted with the same training parameters and random seed used for the networks with 10 and 11 hidden nodes. The results showed that the global-level distinction emerged on 3 of the 8 hidden nodes by 120 sweeps, a point at which the network had not yet successfully distinguished any of the basic-level categories. This finding, replicated with 2 other random seeds, indicates that category representations at the global level precede those at the basic level even when the global-level categories are not taught and the number of hidden nodes matches the number of basic-level representations the network must form. Moreover, the finding provides additional support for the idea that with naturally clustered input sets, the learning route to the perceptual basic level passes through the perceptual global level.

TABLE 1
Number of Hidden Nodes Coding for the Global-Level Distinction between Mammals and Furniture in the Networks of Part V

Number of Training Sweeps	Number of hidden nodes coding the global-level category distinction — Total number of hidden nodes in the network: 8	10	11
120	3	8	5
240	3	8	5
480	2	6	5
960	1	4	4
3,600	1	1	4
7,200	1	1	2
43,200	1	1	2

An interesting result that occurred in all three of the simulations reported in this section of the paper is that the number of hidden nodes coding for the global level decreased with increasing exposure to the stimulus patterns. Table 1 displays the number of hidden nodes representing the global-level distinction between mammals and furniture at a number of different points during training. One can observe that the global level is represented from early in the learning sequence, but that the number of hidden nodes coding for the global level declines steadily. The bottom numbers obtained at 43,200 training sweeps should make it clear, however, that the global level does not completely drop out of the overall pattern of representation. What appears to be the case is a gradual transition to more and more of the representational resources being devoted to coding the basic-level. As subjects begin to encounter objects within a domain, their initial representation of those objects will tend toward the global level. Increasing frequency of experience with objects in that domain results in a greater likelihood that those objects will be represented at the basic or even subordinate levels (see also Schyns, 1991, and Tanaka & Taylor, 1991). The hidden representations thus exhibit different patterns of similarity over the course of training; early they reflect only global-level categories, but later they reflect both global-level and basic-level categories. That is to say, both types of similarity eventually come to coexist in the same representation.

GENERAL DISCUSSION

Connectionist accounts of cognitive development are increasing in number (Elman et al., 1996; Mareschal et al., 1995; McClelland, 1989; Munakata et al., in press; Plunkett & Sinha, 1992; Schyns, 1991). This paper represents one of the first attempts to apply a connectionist analysis to the issue of how perceptual category representations may arise at basic and global levels during early development. A series of network simulations were found to learn categories at basic and global levels and in a global-to-basic sequence. Subsequent simulations revealed that the global-to-basic order was critically dependent on (1) the nature of the global categories (i.e., perceptual vs arbitrary) and (2) whether the network was trained to form perceptual basic-level categories. The order in which the categories emerged may thus be viewed as a consequence of the interaction between a particular ''external'' environment and a specific ''internal'' network architecture.

A striking result of the simulation reported in Part IV 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 may remind some readers of the views of Mandler (in press, 1997) who has argued for distinct perceptual and conceptual levels of category representation in human infants. However, a key difference between our model and the one proposed by Mandler is that in our model a single network (and hence a single system of representation) forms both types of categories, whereas in the Mandler view perceptual and conceptual representations are the products of two complementary, but distinct processes (Mandler & McDonough, 1993). What remains unclear in the present simulations is the precise relation of perceptual and arbitrary (conceptual-like) global-level categories. For example, is the arbitrary global level derived from the perceptual global level in accord with a continuity-based model in which an initial perceptually based representation evolves so as to assimilate more conceptual-like components (e.g., Quinn & Eimas, 1996b)? The fact that the ''global crossover'' network maintained a representation of the perceptual global level in the process of constructing the arbitrary global level is consistent with this view. Alternatively, could a conceptual-like global level be formed by a mechanism that is separate from the one used to form the perceptual global level (cf. Mandler, 1997)? The finding that a model taught to assign stimulus patterns only to arbitrary global-level categories (and not basic-level categories) did not form perceptual global-level categories supports this position. Additional computational and experimental work will clearly be needed to determine what may turn out to be a complex set of relations between perceptual and arbitrary (conceptual-like) global-level categories.

Hypotheses for Experiments

Models are often constructed to fit existing behavioral data. The models reported here have performed well in this respect. Their output was broadly consistent with findings that young infants form both global-level and basic-level category representations (e.g., Quinn & Eimas, 1996b), and also with more recent evidence of Rakison (1996; Rakison & Butterworth, in press) suggesting that infants form global-level category representations on the basis of salient attribute differences. Models are, however, also judged by the degree to which they can generate experimental hypotheses. We therefore offer the following hypotheses for future empirical study:

1. The global-to-basic sequence of category emergence observed in the initial simulation reported in Part I indicates that global-level representations may precede basic-level representations during the course of a fixed exposure period (e.g., single familiarization session) in which exemplars from several basic-level categories from the same global-level category are presented. This hypothesis can be tested with 3- to 4-month-olds, an age group that has already demonstrated the ability to form both global-level and basic-level category representations, albeit in separate experimental sessions (Behl-Chadha, 1996; Quinn et al., 1993).

2. If infants display the ability to form global-level and basic-level representations in a global-to-basic order under one set of task conditions during a single familiarization session, it may 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 an infant is exposed to a stimulus decreases its memorability (Fagan, 1974; Cornell, 1979), 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 category distinctions would be the first to be affected by a moderate decrease in study time per exemplar. Further reduction in study time might affect global-level distinctions as well, but the clear implication is that the basic level would be affected before the global level. These proposed experimental outcomes rest on the assumption that the early-appearing representations that emerge for global-level categories will be more robust than later-appearing representations for basic-level categories (cf. Munakata et al., in press). This assumption is supported in the current simulations by the finding that networks with multiple hidden nodes come to represent the global level from early in the training sequence and continue to represent that level even when the bulk of the representational resources have shifted to the basic level.

3. The multiple hidden node simulations (those conducted with 10 or 11 hidden nodes) indicated that in the course of extended 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. The findings imply that if infants could be repeatedly familiarized with instances of a given category on successive sessions, then there may be a steady transition from global-level to basic-level representation. The simulations also suggest that real-world entities that infants experience on a frequent basis may tend to elicit basic-level responding, albeit subsequent to global-level responding (for corroborating evi-

dence, see Mandler & McDonough, 1993; see also reports of basic-level superiority in older subjects, Horton & Markman, 1980; Mervis & Crisafi, 1982; Rosch et al., 1976).

4. Both global-level and basic-level category representations have been observed with 3- to 4-month-olds in the familiarization–novelty preference procedure (e.g., Quinn & Eimas, 1996b). The global-to-basic sequence observed in the models would therefore suggest that global-level representations should emerge before those at the basic level sometime prior to 3 months of age. This idea can be tested with infants in the age range between birth and 10 weeks that are administered the familiarization–novelty preference procedure.

5. A connectionist approach makes it possible to train models with one or more lesioned input nodes and examine which, if any, category representations fail to emerge. Such manipulations can be helpful in understanding aspects of the input that may be critical for certain category distinctions. For example, the No Face–No Tail model reported in Part II implies that face and tail information may not be necessary for making the category distinction between mammals and furniture, a proposal that can be tested on infants with simple alterations to the mammal stimuli. However, the rate of learning of the global level was slower when the face and tail information was withheld, suggesting that infants' distinction of furniture and ''altered'' mammal stimuli may be more difficult and may require extra familiarization time. The role of salient attribute differences in global-level category differentiation could be further tested by examining the impact of attribute alterations on other global-level category contrasts, e.g., furniture vs vehicles (cf. Rakison, 1996; Rakison & Butterworth, in press).

6. The models presented in Part IV indicate that it may be possible to train subjects, either infants or toddlers, to assign stimuli to arbitrary global-level categories. If the training proceeds by also having the subjects classify the stimuli into basic-level categories, then the simulations suggest that a perceptual global level of category representation will precede formation of both arbitrary global-level categories and basic-level categories. However, if the training on the arbitrary global level occurs without basic-level classification, then the simulations suggest that a perceptual global level of category representation should not be formed. An operant headturning paradigm might be a viable training procedure by which to test these proposals (e.g., Husaim & Cohen, 1981; Kuhl, 1979).

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 (reviewed in Quinn, in press; Quinn & Eimas, 1996b). Data from the experimental studies were used in decisions about what inputs to present to the models. Specifically, the findings that young infants appear to use both external contour and internal feature information from the head and face region to categorically distinguish between cats and dogs led us to assign a number of inputs to attributes from this region of the stimuli (Quinn & Eimas, 1996a; Quinn & Eimas, 1996c; see also, Johnson & Morton, 1991). There is also evidence that infants use correlated attribute information to perform successfully in various kinds of categorization tasks (Younger, 1990, 1992)—a manner of information processing that is broadly consistent with the way in which neural networks learn information.

A limitation of the present approach is that the visual input representations infants (or adults) use to recognize objects are still unknown (see Husaim & Cohen, 1981, and Kemler, 1981, for contrasting views on this issue). While our input scheme did correspond with an object parsing scheme advocated in one contemporary model of object recognition (Zhu & Yuille, 1996), it becomes important to examine whether implementations of our models with a range of input descriptions (e.g., Biederman, 1987; Marr, 1981) would produce comparable results. In our view, incorporating additional features into the input scheme (such as movement, texture, affect, and sound) that might potentially be available to infants for purposes of categorization, would serve to strengthen further the global-level superiority observed during initial category learning. We suspect this would be so because of large differences in attribute values at the global level relative to the basic level for the features mentioned. We therefore believe that our basic observations on perceptual category formation will be robust for the reason that all plausible 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 of the natural and artifactual objects in the environment that we believe to be important for the results obtained, rather than the details of what elements of the visual array are encoded in the input scheme presented to the networks.

A second 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 all of the networks. One can claim that this manner of learning is questionable in the present context for at least two reasons. First, there are many who would 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 about this point. First, at least one level of category representation, the perceptual global level, was obtained without training (see simulations reported in Parts III, IV and V). Second, Plunkett (1996) has noted that processes of cell communication at the level of the synapse are still poorly understood and 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 (Plunkett, 1996). Networks trained on backpropagation commonly develop the same representations as those produced by more biologically plausible, Hebbian learning algorithms. For example, Plaut and Shallice (1993) lesioned a connectionist network trained with a contrastive Hebbian learning algorithm and compared its activity to a lesioned network initially trained with backpropagation. Both networks produced essentially the same pattern of findings.

A third limitation of the models as presently reported is that we have not attempted to model changes in brain structure, peripheral visual capacities (i.e., contrast sensitivity, resolution acuity), and motor coordination, etc., that may occur maturationally during early development. This is not to say that connectionist modeling cannot simulate such changes through alterations to network architecture and input structure (e.g., Elman, 1993; Oliver, Johnson, & Shrager, 1996; Schultz, 1991). While such changes undoubtedly occur in the developing infant, in the present simulations we attempted to examine the range of phenomena that could be accounted for without such alterations.

In conclusion, it will be important to extend many of the effects we have observed in our simulations to other connectionist architectures, input formats, and learning rules (including unsupervised networks). We note that some further issues related to the development of categorization will most likely require simulations with Hebbian self-organizing networks. For example, the question of whether a conceptual global-level category representation can be derived directly from a perceptual global-level category representation may need to be resolved in this way. Despite these acknowledged limitations, we 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 modeling. Such a program may hold promise for the eventual realization of a quantitative and formalized account of category formation by infants.

REFERENCES

Barsalou, L. W. (1983). Ad hoc categories. *Memory & Cognition,* **11,** 221–227.

Behl-Chadha, G. (1996). Basic-level and superordinate-like categorical representations in early infancy. *Cognition,* **60,** 105–141.

Behl-Chadha, G., Eimas, P. D., & Quinn, P. C. (1995, March). *Perceptually-driven superordinate categorization by young infants.* Paper presented at the biennial meeting of the Society for Research in Child Development, Indianapolis, IN.

Biederman, I. (1987). Recognition by components: A theory of human image understanding. *Psychological Review,* **94,** 115–147.

Clark, A. (1993). *Associative engines: Connectionism, concepts, and representational change.* Cambridge, MA: MIT Press.

Clark, A., & Karmiloff-Smith, A. (1993). The cognizer's innards: A psychological and philosophical perspective on the development of thought. *Mind & Language,* **8,** 487–519.

Cornell, E. H. (1979). Infants' recognition memory, forgetting, and savings. *Journal of Experimental Child Psychology,* **28,** 359–374.

Crick, F. H. C. (1989). The real excitement about neural networks. *Nature,* **337,** 129–132.

Eimas, P. D., & Quinn, P. C. (1994). Studies on the formation of perceptually based basic-level categories in young infants. *Child Development,* **65,** 903–917.

Eimas, P. D., Quinn, P. C., & Cowan, P. (1994). Development of exclusivity in perceptually based categories of young infants. *Journal of Experimental Child Psychology,* **58,** 418–431.

Elman, J. L. (1993). Learning and development in neural networks: The importance of starting small. *Cognition,* **48,** 71–99.

Elman, J. L., Bates, E. A., Johnson, M. H., Karmiloff-Smith, A., Parisi, D., & Plunkett, K. (1996). *Rethinking innateness: A connectionist perspective on development.* 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.

Feldman, J. A., & Ballard, D. H. (1982). Connectionist models and their properties. *Cognitive Science,* **6,** 205–254.

Horton, M. S., & Markman, E. M. (1980). Developmental differences in the acquisition of basic and superordinate categories. *Child Development,* **51,** 708–719.

Husaim, J. S., & Cohen, L. B. (1981). Infant learning of ill-defined categories, *Merrill-Palmer Quarterly,* **27,** 443–456.

Johnson, M. H., & Morton, J. (1991). *Biology and cognitive development.* Cambridge, MA: Blackwell.

Karmiloff-Smith, A. (1992a). *Beyond modularity.* Cambridge, MA: MIT Press.

Karmiloff-Smith, A. (1992b). Nature, nurture and PDP: Preposterous developmental postulates? *Connection Science,* **4,** 253–269.

Keil, F. C. (1979). *Semantic and conceptual development: An ontological perspective.* Cambridge, MA: MIT Press.

Kemler, D. (1981). New issues in the study of infant categorization: A reply to Husaim and Cohen. *Merrill-Palmer Quarterly,* **27,** 457–463.

Knapp, A. G., & Anderson, J. A. (1984). Theory of categorization based on distributed memory storage. *Journal of Experimental Psychology: Learning, Memory & Cognition,* **10,** 616–637.

Kohonen, T. (1984). *Self-organization and associative memory.* Berlin: Springer-Verlag.

Kuhl, P. W. (1979). Speech perception in early infancy: Perceptual constancy for spectrally dissimilar vowel categories. *Journal of the Acoustical Society of America,* **66,** 1668–1679.

Mandler, J. M. (in press). Representation. In W. Damon (Ed.), *Handbook of child psychology.* Vol 2: D. Kuhn & R. Siegler (Eds.), *Cognition, perception and language.* New York: Wiley.

Mandler, J. M. (1997). Development of categorization: Perceptual and conceptual categories. In G. Bremner, A. Slater, & G. Butterworth (Eds.), *Infant development: Recent advances* (pp. 163–189). East Sussex, UK: Psychology Press (Erlbaum).

Mandler, J. M., & Bauer, P. (1988). The cradle of categorization: Is the basic level basic? *Cognitive Development,* **3,** 247–264.

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.

Mareschal, D., Plunkett, K., & Harris, P. (1995). Developing object permanence: A connectionist model. In J. D. Moore & J. F. Lehman (Eds.), *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society* (pp. 170–175). Mahwah, NJ: Lawrence Erlbaum Associates.

Marr, D. (1981). *Vision.* San Francisco: Freeman.

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.

McClelland, J. L., & Rumelhart, D. E. (1986a). *Parallel distributed processing: Explorations in the microstructure of cognition.* Vol. 1: *Foundations.* Cambridge, MA: MIT Press.

McClelland, J. L., & Rumelhart, D. E. (1986b). *Parallel distributed processing: Explorations in the microstructure of cognition.* Vol. 2: *Applications.* Cambridge, MA: MIT Press.

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.

Mervis, C. B., & Crisafi, M. A. (1982). Order of acquisition of subordinate, basic, and superordinate categories. *Child Development,* **53,** 258–266.

Miikkulainen, R., & Dyer, M. G. (1991). Natural language processing with modular PDP networks and distributed lexicon. *Cognitive Science,* **15,** 343–399.

Minsky, M., & Papert, S. (1969). *Perceptrons.* Cambridge, MA: MIT Press.

Munakata, Y., McClelland, J. L., Johnson, M. H., & Siegler, R. S. (in press). Rethinking infant knowledge: Toward an adaptive process account of successes and failures in object permanence tasks. *Psychological Review.*

Oliver, A., Johnson, M. H., & Shrager, J. (1996). The emergence of hierarchical clustered representations in a Hebbian neural network model that simulates aspects of development in the neocortex. *Network: Computation in Neural Systems,* **7,** 291–299.

Plaut, D. C., & Shallice, T. (1993). Deep dyslexia: A case study of connectionist neuropsychology. *Cognitive Neuropsychology,* **10,** 377–500.

Plunkett, K., & Elman, J. L. (1997). *Exercises in rethinking innateness: A handbook of connectionist simulations.* Cambridge, MA: MIT Press.

Plunkett, K., & Sinha, C. G. (1992). Connectionism and developmental theory. *British Journal of Developmental Psychology,* **10,** 209–254.

Plunkett, K., Sinha, C. G., Moller, M. F., & Strandsby, O. (1992). Symbol grounding or the emergence of symbols? Vocabulary growth in children and a connectionist net. *Connection Science,* **4,** 293–312.

Plunkett, K. (1996). Development in a connectionist framework: Rethinking the nature-nurture debate. *The Newsletter of the UCSD Center for Research in Language,* **10,** 3–14.

Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Child Psychology,* **77,** 353–363.

Quinn, P. C. (in press). Object and spatial perceptual categorization in young infants: ''What'' and ''where'' in early visual perception. In A. M. Slater (Ed.), *Perceptual development: Visual, auditory and language perception in infancy.* London, UK: UCL Press.

Quinn, P. C., & Eimas, P. D. (1996a). Perceptual cues that permit categorical differentiation of animal species by infants. *Journal of Experimental Child Psychology,* **63,** 189–211.

Quinn, P. C., & Eimas, P. D. (1996b). Perceptual organization and categorization in young infants. In C. Rovee-Collier & L. P. Lipsitt (Eds.), *Advances in infancy research* (Vol. 10, pp. 1–36). Norwood, NJ: Ablex.

Quinn, P. C., & Eimas, P. D. (1996c). Young infants' use of facial information in the categorical differentiation of natural animal species: The effect of inversion. *Infant Behavior & Development,* **19,** 383–386.

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.

Raijmakers, M. E. J., van Koten, S., & Molenaar, P. C. M. (1996). On the validity of simulating stagewise development by means of PDP networks: Application of catastrophe analysis and an experimental test of rule-like network performance. *Cognitive Science,* **20,** 101–136.

Rakison, D. H. (1996, April). *The role of parts and structure in categorization by young children.* Paper presented at the International Conference on Infant Studies, Providence, RI.

Rakison, D. H., & Butterworth, G. E. (in press). Infants' use of object parts in early categorization. *Developmental Psychology.*

Rhodes, G. (1988). Looking at faces: First-order and second-order features as determinants of facial appearance. *Perception,* **17,** 43–63.

Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology,* **7,** 573–605.

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning internal representations by error propagation. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition* (Vol. 1, pp. 318–362). Cambridge, MA: MIT Press.

Rumelhart, D. E., & Todd, P. M. (1993). Learning and connectionist representations. In D. E. Meyer & S. Kornblum (Eds.), *Attention and performance XIV: Synergies in experimental psychology, artificial intelligence, and cognitive neuroscience* (pp. 3–30). Cambridge, MA: MIT Press.

Schultz, T. R. (1991). Simulating stages of human cognitive development with connectionist models. In L. Birnbaum & G. Collins (Eds.), *Machine learning: Proceedings of the Eighth International Workshop* (pp. 105–109). San Mateo, CA: Morgan Kaufman.

Schyns, P. (1991). A modular neural network model of concept acquisition. *Cognitive Science,* **15,** 461–508.

Sherman, T. (1985). Categorization skills in infants. *Child Development,* **56,** 1561–1573.

Smolensky, P. (1988). On the proper treatment of connectionism. *Behavioral and Brain Sciences,* **11,** 1–74.

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.

Tanaka, J. W., & Taylor, M. (1991). Object categorization and expertise: Is the basic level in the eye of the beholder? *Cognitive Psychology,* **23,** 457–482.

Yamane, S., Kaji, S., & Kawano, K. (1988). What facial features activate face neurons in the inferotemporal cortex of the monkey? *Experimental Brain Research,* **73,** 209–214.

Young, M. P., & Yamane, S. (1992). Sparse population coding of faces in inferotemporal cortex. *Science,* **256,** 1327–1331.

Younger, B. A. (1990). Infants' detection of correlations among feature categories. *Child Development,* **61,** 614–620.

Younger, B. A. (1992). Developmental change in infant categorization: The perception of correlations among facial features. *Child Development,* **63,** 1526–1535.

Zhu, S. C., & Yuille, A. L. (1996). FORMS: A flexible object recognition and modeling system. *International Journal of Computer Vision,* **20,** 187–212.

RECEIVED: March 7, 1996; REVISED: April 14, 1997.