Connectionist Models of Infant Perceptual and Cognitive Development

Leslie B. Cohen and Harold H. Chaput

University of Texas at Austin

Commentary submitted to Developmental Science (2001).

Connectionist Models of Infant Perceptual Development

In recent years an exciting innovation has occurred in the study of early development. Connectionist models have begun to appear on a variety of aspects of infant perception, cognition and language. These models have included simulations of object permanence (Munakata, McClelland, Johnson & Sigler, 1997); categorization (Mareschal & French, 2000); causal perception (Chaput & Cohen, in press); early word learning (Schafer and Mareschal, 2001); and simple rule learning (Shultz & Bale, in press). Continuing in this tradition, Mareschal and Johnson now present a connectionist model of a cornerstone in infant visual perception, the perception of object unity. This is an ideal topic to model for a number of reasons: As the authors note, the perception of object unity is subject to a continuing theoretical debate about the role of innate core knowledge versus early experience and learning; there is also a wealth of empirical evidence with infants upon which to draw; and some of that empirical evidence suggests a developmental change in how occluded objects (or object parts) are perceived.

Mareschal and Johnson's model includes several attractive features. It is essentially a constructive model that builds the percept of one versus two objects from a set of lower order perceptual cues. Most, if not all, of these perceptual cues (e.g., the presence of motion, texture, or T junctions) make sense, and they are probably available to young infants. The model learns from environmental experience and generalizes to novel instances; it also produces some, albeit weak, evidence for a developmental change.

Nevertheless, the present model, like many earlier connectionist models, serves mainly as an existence proof, meant to demonstrate that such a system can learn some concept. In that sense it provides a logical counter-argument to postulations of certain types of innate core knowledge. But the field has now advanced to the point where it is important to go beyond existence proofs and attempt to model more closely how these concepts are actually learned by infants and how developmental changes in these concepts actually occur.

We too are working on a connectionist model of infant cognitive development (Chaput and Cohen, in press), and we would be the first to admit that our own attempts influence our view of other models. Nevertheless, from this experience we would like to raise a few general issues we have confronted that also apply to the Mareschal and Johnson model. These are 1) choice of a connectionist architecture, 2) modeling habituation vs. long-term experience, 3) the realism of the training environment, and 4) modeling developmental change.

A growing trend in connectionist models of infant development is the use of architectures that are unsupervised and self-organizing. We agree with this trend, as it avoids many of the problems associated with the use of an external "target" against which the model's output is compared. It is unclear where, in infants, a target would come from, or how such a target would be chosen in a consistent way throughout development and across domains. Auto-associative networks answer this question quite nicely by making the target identical to the input. In our own work, we avoid the question entirely by using Self-Organizing Maps (Kohonen, 1997), which make no use of targets and are, thus, unsupervised. Although Mareschal and Johnson claim that their model is unsupervised, their use of an external target, via the "Direct Route," makes their model supervised by definition. And the source of this target is not actually the input vector (as it would be in an auto-associative network) but the result of an algorithmic manipulation of environmental features not present in the input vector. So their model is not really self-organizing, either. In spite of these possible criticisms, to us, the specific architecture of the model is less important that the results it produces.

Most infant experiments employ a habituation technique to explore perceptual or cognitive development. This includes the research on object unity cited by Mareschal and Johnson; research they presumably are trying to model. But habituation can be used either to assess short-term learning <u>or</u> to probe existing knowledge structures that are the result of long-term experience. We believe it is important for the model to be explicit about these two types of learning, as our model does. The distinction is unclear in the Mareschal and Johnson model. If their training is meant to represent short-term habituation, then it is unlike what infants have received in actual experiments since it includes exposure to an unoccluded event in addition to an occluded event. On the other hand if the training is meant to represent long-term experience, then it is unclear how habituation is represented in their model at all.

One can also ask how realistic or artificial that long-term training is. We believe, as Mareschal and Johnson do, that infants are able to learn concepts by processing information they receive from the environment. But we have doubts about the plausibility of the environment provided in Mareschal and Johnson's models, especially Model 2. The authors posit that an infant's bias to perceive a partially occluded rod as two separate rods is the result of an "inductive bias" which is present, not only in 4-month-olds, but also in neonates. They go on to argue that this "inductive bias" comes about "as a result of experience with a particular environment.” Earlier, though, they note that Model 2 prefers the two-rod response “because most of the events in the learning environment arise from two objects.” So we wonder if they might be “giving away the answer” by skewing the environment in such a way as to facilitate a certain outcome. Our intuition is that an infant's environment is more likely to contain whole rods than co-moving, parallel, relatable rod segments.

Finally, Mareschal and Johnson’s model is only one part of a much larger story of cognitive development, and we would be interested in seeing if their model can cover more of this story. Specifically, with regard to object unity, Eizenman and Bertenthal (1998) performed a similar bar-and-occlusion study with infants; only instead of moving the bar laterally behind the occluder, they rotated the bar. Rotation made the perceptual job more difficult for the infant, and they found that 4-month-olds tended to regress to a two-bar bias, whereas 6-month-olds finally regained the single-bar bias. Their results suggest the presence of an additional transition in the development of object unity, and we are curious to see if Mareschal and Johnson’s model can capture this transition as well.

We also believe the set of developmental transitions found in infants' perception of object unity represent but a single example of a more general set of principles that appear to be at work across many domains throughout early cognitive development, (Cohen, 1998; Cohen & Cashon, 2001). These principles provide a constructivist view of cognitive development amenable to connectionist modeling. Our own work (Chaput & Cohen, in press) is a beginning attempt to address this broader picture with a single connectionist architecture that adheres to these principles.

We hope these comments are taken in the positive way they are intended. We believe the authors' approach, along with the other approaches we have mentioned, such as auto-associative networks, cascade correlation models, or our own hierarchically arranged self-organizing maps, hold considerable promise for explaining the development of early perceptual and cognitive ability. We are pleased to see that connectionist models are progressing beyond existence proofs and attempting more complete and potentially compelling explanations. In the end these models will be evaluated by more than their ability to simulate existing empirical data. They will also be required to generate new developmental predictions that are then confirmed empirically. Once they reach that stage they clearly will be making a major contribution to our understanding of early perceptual and cognitive development.

References

Chaput, H. H. & Cohen, L. B. (In press). A model of infant causal perception and its development. Proceedings of the 2001 Cognitive Science Society Meeting, Edinburgh.

Cohen, L.B. (1998). An information-processing approach to infant perception and cognition. In F. Simion & G. Butterworth (Eds.), The Development of Sensory, Motor, and Cognitive Capacities in Early Infancy (pp.277-300). East Sussex: Psychology Press.

Cohen, L. B. & Cashon, C. H. (2001). Infant object segregation implies information integration. Journal of Experimental Child Psychology. 78, 75-83.

Eizenman, D. R. & Bertenthal, B. I. (1998). Infants' perception of object unity in translating and rotating displays. Developmental Psychology, 34, 426-434.

Kohonen, T. (1997). Self-Organizing Maps. Berlin: Springer-Verlag.

Mareschal, D. & French, R. M. (2000). Mechanisms of categorization in infancy. Infancy, 1, 59-76.

Munakata, Y., McClelland, J. L., Johnson, M. H., & Siegler, R. S. (1997). Rethinking infant knowledge: Toward an adaptive process account of successes and failures in object permanence tasks. Psychological Review, 104, 686-713.

Shafer, G. & Mareschal, D. (2001). Modeling infant speech sound discrimination using simple associate networks. Infancy, 2, 7-28.

Shultz, T. R. & Bale, A. C. (In press). Neural network simulation of infant familiarization to artificial sentences; Rule-like behavior without explicit rules and variables. Infancy.

Author Notes

Leslie B. Cohen, Department of Psychology

Harold H. Chaput, Department of Computer Science

Preparation of this paper was supported in part by NIH grant HD-23397 to the first author from the National Institute of Child Health and Human Development. Written correspondence and requests for reprints should be sent to Leslie B. Cohen, Department of Psychology, Mezes Hall 330, University of Texas, Austin, TX 78712. Electronic correspondence may be sent via Internet to [cohen@psy.utexas.edu] or to [chaput@cs.utexas.edu].