Babies, Variables, and Connectionist Networks

Michael Gasser (GASSER@CS.INDIANA.EDU)
Eliana Colunga (ECOLUNGA@CS.INDIANA.EDU)
Computer Science Department, Cognitive Science Program
Indiana University
Bloomington, IN 47405

Two recent papers in *Science* have demonstrated the remarkable language learning abilities that are possessed by infants. In both cases the infants were presented with sequences of syllables embodying some sort of regularity and later tested with sequences that agreed or disagreed in certain ways with the training set. In the experiments of Saffran, Aslin, and Newport (1996), eight-month-olds were able to distinguish three-syllable "words" that they had heard previously in a stream of syllables from those that they had not. In the experiments of Marcus, Vijayan, Bandi Rao, and Vishton (1999), seven-month-olds were able to distinguish three-syllable "sentences" that followed a training pattern, either AAB, ABB, or ABA (where A and B are variables representing syllables), from those that did not. What is striking about the latter experiments is that the test sentences consisted of syllables not heard at all during training. That is, the pattern that is learned seems to be independent of specific syllable content and to require variables for its representation.

Marcus et al. (1999) and Pinker (1999) argue that these experiments demonstrate two distinct learning mechanisms, which will later prove useful for learning language: a statistical/associationist mechanism, illustrated in Saffran et al.'s experiments, and a symbolic mechanism, illustrated in Marcus et al.'s experiments. In particular, because variables are excluded from simple associationist accounts, these accounts fail to handle Marcus et al.'s results.

We focus on Marcus et al.'s experiments, showing how an associationist device, a particular neural network architecture, can learn the patterns in the experiments, generalizing to novel sequences, and how this account, rather than being simply an uninteresting implementation of a symbolic model, makes novel predictions about the learning of sequences.

A neural network needs two features to generalize over "grammatical" patterns like those in Marcus et al.'s experiments. First, it needs a means of distinguishing particular objects or sets of objects from each other in short-term memory. This allows it to *segment* input sequences on the basis of inter-item similarity, for example, for the sequence *le le di*, to treat the first and second syllables as belonging to a group distinct from the third. Second, it needs a means of handling relational knowledge. This allows it to detect and store the similarity relations that characterize segmented sequences, for example, to remember that the first and second syllables in a sequence resembled each other but differed from the third.

Playpen (Gasser & Colunga, 1998) is a neural network architecture which is designed to represent and learn relational knowledge and to deal with simple sequential patterns. Segmentation is handled through the use a second dimension, in addition to activation, along which units in the network can vary. Units which are synchronized along this dimension are treated as belonging together. Relational knowledge in Playpen is handled with special-purpose *micro-relation units* (MRUs), which are activated to the extent that input to their two *micro-roles* is out of synchronization, that is, that it represents distinct objects or groups.

To simulate Marcus et al.'s task, we used an instantiation of Playpen in which similarity relations between the syllable input units were represented by hard-wired connections. These connections cause similar syllables in a sequence to be synchronized and different syllables to be desynchronized. MRUs in the network represented possible binary relations between the syllables in a sequence, and connections between the MRUs represented correlations between these simple binary relations.

The network was trained on sequences of syllables adhering to one or another of the patterns used in Marcus et al.'s experiments, and it successfully generalized to sequences of novel syllables by producing more activation across the MRU layer for "grammatical" than for "ungrammatical" sequences. But in the model this rule-like behavior is just one extreme of a continuum along which sensitivity to item content varies. Specifically, in the network the similarity of the test sequences to the training sequences had an effect on generalization. This behavior contrasts with that of symbolic accounts, where the variables in the learned rule would be oblivious to the specific content of the items in the test patterns.

References

Gasser, M. & Colunga, E. (1998). Where do relations come from?. Tech. rep. 221, Indiana University, Cognitive Science Program, Bloomington, IN.

Marcus, G. F., Vijayan, S., Bandi Rao, S., & Vishton, P. M. (1999). Rule learning by seven-month-old infants. *Science, 283*, 77–80.

Pinker, S. (1999). Out of the minds of babes. *Science, 283*, 40–41.

Saffran, J., Aslin, R., & Newport, E. (1996). Statistical learning by eight-month-old infants. *Science, 274*, 1926–1928.