The Role of Frequency in Modeling Double Dissociations

Laura M. Gonnerman (LAURA@GIZMO.USC.EDU)
Michael W. Harm (MHARM@GIZMO.USC.EDU)
Elaine S. Andersen (ELAINE@GIZMO.USC.EDU)
Program in Neural, Informational, and Behavioral Sciences
University of Southern California, Los Angeles CA 90089-2520

Introduction

Connectionist models have been used to demonstrate that double dissociations can arise in non-modular systems. In a recent paper, Bullinaria and Chater (1995) argue against this general approach, questioning the reliability of double dissociations produced in non-modular systems. Focusing on double dissociations based on rules versus exception cases, they create simulations which show that in large-scale, fully distributed networks, the rule based exemplars are retained, and the exception cases are lost. The opposite dissociation, where the exceptions are spared and the rules are lost, is *only* obtained in small-scale, less biologically plausible networks. Based on these results, they challenge the logic of using connectionist simulations to refute the modularity hypothesis.

Bullinaria and Chater's simulations are based on artificial training example sets. While they manipulated regularity, having some examples that were rule governed, and others that were exceptions, they failed to manipulate the relative frequencies of the training exemplars.

The frequency of different tokens and types in natural language has been explored in several domains. For example, children acquire more frequent forms earlier than less frequent ones (e.g. Clark & Berman, 1984), and exceptions that are low in frequency tend to be regularized over time through diachronic processes of language change (Haiman, 1994). In addition, there is an interaction observed between frequency and regularity, such that high frequency items are less affected by irregularity (Seidenberg et al., 1984). Therefore, the frequency of an irregular item is fundamental to any consideration of its behavior in a language system.

We hypothesized that Bullinaria and Chater failed to get a double dissociation because their training set did not reflect the frequency properties of natural language. Working within their paradigm, we explore this hypothesis, examining both the frequency of different tokens and the distribution of types in the simulations described below.

The Models

The models described here were the same as those in Bullinaria and Chater (1995), except that the frequency of the training exemplars was manipulated. Two training sets were created based on the stimuli used by Bullinaria and Chater: one was comprised of rule and sub-rule exemplars; the other consisted of rules and exceptions. The English past tense is an example of a language system that incorporates rule-governed (*bake-baked*), sub-rule (*ring-rang, sing-sang*), and exception (*go-went*) cases. Therefore, training exemplars were assigned frequency values according to the actual frequency distribution of the English past tense. In training, we weighted the error signal according to the raw frequency value assigned to each exemplar. For the first set of simulations, the hidden unit layer included 100 units, while for the larger models, 600 hidden units were used. The models were damaged in two ways, by removing either hidden units or connections.

Results and Discussion

As expected, the small networks with the frequency manipulation demonstrated double dissociations; in some instances performance on the rule-governed items was worse than performance on the sub-rule or exception items, and in other cases the reverse pattern occurred. However, in contrast to Bullinaria and Chater's simulations, our large networks also produced the double dissociation, such that on relatively rare occasions (as with patients), the sub-rules or exceptions were spared compared to the rules.

The simulations reported here show that frequency manipulations are crucial for creating robust networks that learn sub-rules as well as true exceptions. It is a well established fact in natural language that low frequency exceptions are regularized by children, lost in brain damage, and regularized over time through diachronic processes of language change.

In light of the results presented here, the challenge Bullinaria and Chater set forth to connectionist researchers may well be met if considerations of the nature of the input, in this instance the frequency of irregular and exception cases, is taken into account in simulating natural language and brain damage.

References

Bullinaria, J., & Chater, N. (1995). Connectionist modelling: Implications for cognitive neuropsychology. *Language and Cognitive Processes, 10*, 227-264.

Clark, E., & Berman, R. (1984). Structure and use in the acquisition of word formation. *Language, 60*, 542-590.

Haiman, J. (1994). Ritualization and the development of language. In W. Pagliuca (Ed.), *Perspectives on grammaticalization*. Amsterdam: John Benjamins.

Seidenberg, M. S., Waters, G. S., Sanders, M., & Langer, P. (1984). Pre- and postlexical loci of contextual effects on word recognition. *Memory and Cognition, 12*(4), 315-328.