

Question Formation, Neural Networks, and the Poverty of the Stimulus

Robert Frank¹, Donald Mathis¹, Ebonye Gussine¹, John Stowe¹, and Manny Vindiola¹

¹Department of Cognitive Science, Johns Hopkins University

{rfrank, mathis, ebonye, jstowe, mannyv}@jhu.edu

Chomsky's (1975) argument from the poverty of the stimulus (APS) has recently been disputed by Lewis and Elman (2001, LE), who suggest that even in the absence of crucial data like (i) a Simple Recurrent Network (SRN) can induce a structure sensitive generalization for question formation. LE trained an SRN to perform word prediction on declarative and interrogative English sentences, though not examples like (i).

(i) *Is the bird that is singing lonely?*

They demonstrate that the network generalizes in an apparently structure sensitive fashion when tested on cases like (i): at the relative pronoun 'that' the network predicts an auxiliary verb, while at the end of the relative clause it fails to predict an auxiliary.

There is cause for skepticism, however. First, LE's network does not encode any relationship between declarative and interrogative forms, as this is unnecessary for the prediction task. Yet since structure sensitivity arises only in the context of this relation, the simulation simply doesn't bear on whether an innate structural bias is necessary. Secondly, Reali and Christiansen (2005) show that the distinction between (i) and its non-structure sensitive counterpart (ii) can be predicted through non-structural means, using a bigram language model.

(ii) *Is the bird that singing is lonely?*

To assess structural knowledge more directly, we trained SRNs to map between sentence types. Adapting the set-up explored by Botvinick and Plaut (in press), we presented SRNs with declarative sentences one word at a time, followed by a recall cue. During sentence presentation, the network's target outputs were those of word prediction. However, after the recall cue the network was trained to reproduce either the original declarative sentence (when a D cue was given) or a question variant (Q cue), one word at a time. The training set of 100K examples included sentences with D and Q triggers in equal numbers, with and without auxiliary verbs, transitive and intransitive verbs, singular and plural nouns and verbs (with agreement), and PP and relative clause modifiers. However, sentences like (i) were withheld for Q-triggers. The network performed quite successfully, producing both question and declarative versions of input *for sentence types in which it received training*. Errors consisted almost entirely of substitutions of grammatically appropriate words. However, when a Q-trigger was presented following a sentence with subject modified by a relative clause, the network failed, producing a question of a different sentence type. Our results suggest that the network is incapable of generalizing a single structure-sensitive rule of question formation across different structures, instead forming distinctive mappings for each sentence type. We take this to undercut claims that SRNs provide a means of overcoming the APS.

References

- Botvinick, M. and Plaut, D. C. (in press). Short-term memory for serial order: A recurrent neural network model. *Psychological Review*.
- Chomsky, N. (1975). *Reflections on Language*. New York: Pantheon.
- Lewis, J. D., and Elman, J. L. (2001). Learnability and the statistical structure of language: Poverty of stimulus arguments revisited. In *Proceedings of the 26th Annual Boston University Conference on Language Development*.
- Reali, F., and Christiansen, M. H. (2005). Uncovering the richness of the stimulus: Structure dependence and statistical evidence. *Cognitive Science* 29, 1007-1028.