

Connectionist Semantic Systematicity in Language Production

A defining characteristic of human language is systematicity: “the ability to produce/understand some sentences is intrinsically connected to the ability to produce/understand certain others” (Fodor & Pylyshyn, 1988). Further, Fodor and Pylyshyn (1988) argue that connectionist models are not able to display systematicity without implementing a classical symbol system.

The connectionist comprehension model developed by Frank, Haselager, and van Rooij (2009), however, challenges this highly debated assertion, by developing a connectionist model of comprehension which is argued to achieve relevant levels of systematicity. Their model constructs a situation model (see Zwaan and Radvansky (1998)) of the state-of-affairs described by a sentence that also incorporates world knowledge-driven inferences. When the model processes a sentence like ‘a boy plays soccer’, for instance, it not only recovers the explicit, literal propositional content, but also constructs a more complete situation model in which a boy is likely playing outside on a field, with a ball, with others, and so forth. Crucially, Frank et al. (2009)’s model generalizes to both sentences and situations that it has not seen during training, exhibiting different levels of semantic systematicity and is argued to provide an important step in the direction of psychologically plausible models of language comprehension.

In the present paper, we examine whether the approach developed by Frank et al. (2009) is equally well suited to language production, and present a connectionist production model that generates sentences from these situation models.

We employ an extended Simple Recurrent Neural Network architecture (SRN) (Elman, 1990). Our architecture is broadly similar to the one used by Frank et al. (2009), with the main difference being that the inputs and outputs are reversed; it maps situation model representations onto sequences of localist word representations.

In order to assess the performance of the model, we tested it on 5 different conditions representing different levels of generalization or systematicity. In all cases, the queried sentence type has never been seen by the model. We defined a similarity score to evaluate the results. On the training set, the model achieved an average similarity score of 99.43% (and 98.23% perfect matches). On the testing set, the average similarity score across all conditions is of 97.1%, with 88.57% of perfect matches.

Although the performance of the model is very high, the model elicits some mistakes that allow us to get some insight into the internal mechanism of the model. After an analysis of the output, we see that all the sentences produced are syntactically correct and semantically felicitous. The vast majority of the elicited mistakes occur when the model produces a sentence that is semantically highly similar to the one expected. We hypothesize that this happens because the model is able to roughly reconstruct the semantic space, putting together representations that are semantically similar and thus assigning similar linguistic representations to them.

We conclude that our model successfully learns to produce sentences from the situation models. Importantly, we demonstrate that this model is able to describe both unseen situations, demonstrating semantic systematicity similar to Frank et al. (2009), as well as produce alternative encodings (e.g. active/passive) for a given situation, that were not seen during training and thus demonstrating syntactic systematicity.

Primary author: Mr CALVILLO, Jesus (Saarland University)

Co-authors: Dr BROUWER, Harm (Saarland University); Prof. CROCKER, Matthew (Saarland University)

Presenter: Mr CALVILLO, Jesus (Saarland University)

Track Classification: Student