Tel Aviv University
School of Philosophy, Linguistics and Science Studies,
Department of Linguistics

THURSDAY INTERDISCIPLINARY COLLOQUIUM

Thursday 19/06/2025

16:15-17:45
via Zoom

Matan Abudy, Tel Aviv University

Symbolic Generalization in Neural Networks with Minimum Description Length

Language models and neural networks may seem flawless, but are they truly perfect? A closer look reveals that they often settle for approximations rather than exact solutions. Why?

We step back and ask: What principles guide language learning? How does a child generalize rules from limited evidence? It has been suggested that children, like scientists, propose theories to explain observed data and choose the simplest one that fits. This idea is formally known as the Minimum Description Length (MDL) principle.

In recent years, MDL has been successfully applied to a wide range of linguistic phenomena, where generalization from limited data is crucial. In phonology, learning models based on the MDL principle are currently the only ones that infer linguistic knowledge from distributional evidence - even in challenging cases involving opacity and optionality.

MDL has also been applied to neural networks. Previous work showed that applying MDL to Recurrent Neural Networks (RNNs) enabled them to learn artificial grammars perfectly from scratch, something they couldn’t do otherwise.

We evaluate MDL on two new grammars: one modeling arithmetic exercises, the other a fragment of English. We find that a key component of how neural networks are trained to solve tasks—regularization methods—plays a crucial role. Standard methods not only fail to find perfect solutions but actively push networks away from them.

We compare standard methods to MDL and validate its robustness across different learning setups. We propose that MDL introduces the appropriate inductive bias for true symbolic generalization.

Click here to see the colloquium program.