Large language models and human cognition

Dr ROGER LEVY

Department of Brain and Cognitive Sciences at the Massachusetts Institute of Technology
- MIT's Computational Psycholinguistics Laboratory

Large language models and human cognition Large language models (LLMs) are the first human-created artifacts whose text processing and generation capabilities seem to approach our own. But the hardware they run on is vastly different than ours, and the software implementing them probably is too. How, then, can we use LLMs to advance the science of language in the human mind? In this talk I present a set of case studies that exemplify three answers to this question: LLMs can help us place lower bounds on the learnability of linguistic generalizations; they can help us reverse-engineer human language processing mechanisms; and they can help us develop hypotheses for the interface between language and other cognitive mechanisms. The case studies include controlled tests of grammatical generalizations in LLMs; computational models of how adults understand what young children say; psychometric benchmarking of multimodal LLMs; and neurosymbolic models of reasoning in logical problems posed in natural language.

This talk covers joint work with Elika Bergelson, Veronica Boyce, Ruth Foushee, Richard Futrell, Alex Gu, Jennifer Hu, Anna Ivanova, Benjamin Lipkin, Gary Lupyan, Kyle Mahowald, Stephan Meylan, Theo Olausson, Subha Nawer Pushpita, Armando Solar-Lezama, Joshua Tenenbaum, Pranali Vani, Ethan Wilcox, Nicole Wong, and Cedegao Zhang.