Architecture- and data-constrained LLMs don't produce childlike language

Dyana Muller
Department of Neuroscience, UC Berkeley
Steven Piantadosi
Department of Psychology, UC Berkeley
Department of Neuroscience, UC Berkeley

Abstract

Large language models (LLMs) demonstrate impressive, adult-level language capabilities. State-of-the-art models produce syntactically correct, coherent, and fluent text that is often indistinguishable from human writing. However, their validity as models of language acquisition is less clear. A compelling model of the acquisition process should not only capture mature language use but also reproduce the characteristic syntactic and semantic errors children make during learning. To test this, we simulate children's cognitive and experiential constraints by limiting GPT-2 architecture and training data, and then compare the language output to child language. Variants of GPT-2 constrained in number of transformer blocks, number of attention heads, embedding dimensionality, context length, or number of training steps all fail to exhibit childlike linguistic patterns. These findings suggest that current LLMs are poor theories of the human language acquisition process.

In G. Boleda, I. Dautriche, M. Maldonado, C. Saldana and T. Torrent (Eds.), *Proceedings of the 48th Annual Conference of the Cognitive Science Society*