

Large Language Models and Human Discourse Processing

Eyal Sagi

University of St. Francis, Joliet, Illinois, United States

Abstract

Recent advances in generative language models, such as ChatGPT, have demonstrated an uncanny ability to produce texts that appear to be comparable to those produced by humans. Several key empirical results related to human processing of language, such as analogical reasoning, have been replicated using these models. Nevertheless, there are some important differences between the language generated by these models and language produced by humans. In this paper, I examine how LLMs performs on two pronoun disambiguation tasks reported on by Rhode, Levy, and Kehler (2011) and Sagi and Rips (2014). While LLMs performed reasonably in these tasks, their responses demonstrate stronger language-based biases while the influence of world knowledge, such as causal relationships, was lessened. Because LLMs replicate language produced by humans, these results can help shed light on which aspects of language use are directly encoded in language and which require additional reasoning faculties beyond language processing.