

Linguistic Framing in Large Language Models

Stephen Flusberg

Vassar College, Poughkeepsie, New York, United States

Kevin Holmes

Reed College, Portland, Oregon, United States

Abstract

Large Language Models (LLMs) have captured the world's attention for their surprisingly sophisticated linguistic abilities, but what they might reveal about human cognition remains unclear. Meanwhile, members of the public routinely share "prompt engineering" tips for eliciting "better" responses from LLMs such as OpenAI's ChatGPT. These efforts parallel research on linguistic framing, which shows that subtle linguistic cues shape people's attitudes and decision-making in a variety of contexts. In this study, we tested whether state-of-the-art LLMs would exhibit similar framing effects as human participants. We adapted a range of linguistic framing stimuli for use with LLMs based on a recently developed taxonomy of framing effects (e.g., lexical, figurative, and grammatical framing). Results revealed that some but not all framing effects replicated with LLMs. These findings have practical applications for interacting with AI systems and inform our understanding of the cognitive mechanisms that underlie the effects of framing.