JOIN US

FRIDAY

MARCH

10

12pm (noon) ET

Presentation in Zoom, accessible via the C-STAR website:

http://cstar.sc.edu/lecture-series/

C-Star Lecture Series

THE UTILITY OF LARGE LANGUAGE MODELS IN UNDERSTANDING THE HUMAN LANGUAGE SYSTEM

Evelina Fedorenko, PhD

Massachusetts Institute of Technology

Until a few years ago, the field of language research has lacked computational models that can take as input an arbitrary linguistic string and represent it in a way so as to enable a broad range of downstream tasks, from text generation, to question answering, to translation into another language. Now, we have a class of such models—commonly referred to as 'large language models' (LLMs)—that have taken the fields of Natural Language Processing and Artificial Intelligence by storm. Aside from revolutionizing AI, these models have tremendous potential for transforming the field of language research, including potential applications in work on developmental and acquired language disorders. In this talk, I will start by presenting evidence that representations extracted from some LLMs can capture human neural responses (as measured with fMRI or intracranial recordings) during the processing of the same language stimuli. Moreover, models that perform better on the next-word prediction task are better able to predict human responses, but performance on other language tasks (e.g., grammaticality judgments) does not bear a relationship to brain predictivity. This selective relationship between next-word prediction performance and brain predictivity suggests that optimizing for predictive representations may be a shared objective of both biological and in silico language systems. I will then talk about more recent work that builds on these original findings to better understand the necessary and sufficient conditions for a model to capture human neural responses. I will talk about a study that shows that models that are trained on developmentally plausible amounts of data already provide a good match to human neural data. This result helps address the major criticism of LLMs as models of human language processing: namely, that they are trained on vastly more data than what human children get exposed to. Finally, I will talk about a series of studies where the model representations are altered before relating them to human neural data in an effort to identify the critical features of the representations that mediate the model-to-brain match. These studies suggest that the key contributor to the model-to-brain match is the lexical-semantic content rather than syntactic cues like word order and function words. I will conclude by outlining promising future directions of this research program, including its clinical applications.

UNIVERSITY OF SOUTH CAROLINA
Arnold School of Public Health

The online lecture can be followed online from your computer, tablet, or smartphone in Zoom.
The Zoom link is accessible via the C-STAR website: http://cstar.sc.edu/lecture-series/