International Conference on Machine Learning, Artificial Intelligence and Data Science (ICMLAIDS 2026)

Conference Proceedings   March 20 - 21, 2026 - Virtual

Do Language Models Know Language?

Gerald Penn

Professor of Computer Science, University of Toronto, USA

Abstract

Triumphalist portraits of large language models (LLMs) boast that language models have mastered a level of language understanding that natural language processing (NLP) researchers have laboured for years to attain using complex architectures consisting of diverse component models that each require large amounts of training data.

Do they? How do we know? And, if so, how do they manage this?
In this talk, we will examine some recent results in relation to LLMs that cast doubt upon these claims, while affirming the utility of LLMs in present-day NLP.