

**Cameron Jones:
Large language models, humanlikeness, and
authenticity**

Abstract:

Large Language Models (LLMs) have proven to be strikingly effective at mimicking superficial patterns in human behaviour, raising theoretical questions about the uniqueness of human cognition, and practical questions about the authenticity of text we encounter online. While the Turing test has historically been treated as a measure of machine intelligence, here I will discuss the Turing test as a method for investigating authenticity and humanlikeness. Across 3 experiments we find people struggle to distinguish between humans and LLMs in short conversations, suggesting that these systems pass the Turing test. In this talk I will discuss what these results mean for our conception of humanlikeness and authenticity, and the aspects of our behaviour that are most uniquely and authentically human.