

Tom Griffiths

Department of Psychology, UC Berkeley

Connecting human and machine learning via probabilistic models of cognition

Human performance defines the standard that machine learning systems aspire to in many areas, such as forming new concepts, making scientific discoveries, and learning language. This suggests that studying human cognition may be a good way to develop better learning algorithms, as well as providing basic insights into how the mind works. However, in order for ideas to flow easily from psychology to computer science and vice versa, we need a common language for describing human and machine learning. I will summarize recent work exploring the hypothesis that probabilistic models of cognition, which view learning as a form of statistical inference, provide such a language, including results that illustrate how novel ideas from statistics can inform cognitive science.