Modeling Open-World Cognition as On-Demand Synthesis of Probabilistic Models

Lionel Wong
Stanford University, Palo Alto, California, United States

Katherine Collins
MIT, Cambridge, Massachusetts, United States

Lance Ying
Massachusetts Institute of Technology, Cambridge, Massachusetts, United States

Cedegao Zhang
Massachusetts Institute of Technology, Cambridge, Massachusetts, United States

Adrian Weller
University of Cambridge, Cambridge, United Kingdom

Tobias Gerstenberg
Stanford University, Stanford, California, United States

Timothy O'Donnell
McGill University, Montreal, Quebec, Canada

Alexander Lew
MIT, Cambridge, Massachusetts, United States

Jacob Andreas
MIT, Cambridge, Massachusetts, United States

Tyler Brooke-Wilson
Yale University, New Haven, Connecticut, United States

Joshua Tenenbaum
MIT, Cambridge, Massachusetts, United States

Abstract

People are able to reason flexibly across a vast range of domains and contexts, from navigating new environments and social situations, to playing new games and even betting on the outcomes of new sports. How do we draw on our knowledge and past experiences to tractably make sense of any particular situation? Here, we explore the hypothesis that people use a combination of distributed and structured, symbolic knowledge to construct bespoke mental models tailored to novel situations. We propose a computational implementation of this idea – a 'Model Synthesis Architecture' ('MSA') – using language models as a stand-in for distributional knowledge and a probabilistic programming language to express bespoke probabilistic symbolic models. We evaluate our model with respect to human judgments on a novel reasoning dataset. Our "Model Olympics" domain comprises a series of "sports commentary" vignettes, and is designed to test open-ended reasoning by requiring (i) reasoning about arbitrary causal structures described in language; (ii) drawing in relevant latent considerations from background knowledge; and (iii) flexibly adapting to an 'open world' setting with novel observations sourced from other human participants. We compare our MSA to hand-coded probabilistic programs and LM-only baselines. We find that our approach captures key hallmarks of rational inference from human judgments that the LM-only baselines do not, especially for very novel scenarios. See https://sites.google.com/view/openworld-msa?usp=sharing for additional details and preprint.

In D. Barner, N.R. Bramley, A. Ruggeri and C.M. Walker (Eds.), *Proceedings of the 47th Annual Conference of the Cognitive Science Society*