Concept Learning as Coarse-to-Fine Probabilistic Program Induction

Maddy Bowers
MIT, Cambridge, Massachusetts, United States

Alexander Lew
MIT, Cambridge, Massachusetts, United States

Wenhao Qi
University of California, San Diego, La Jolla, California, United States

Joshua Rule
UC Berkeley, Berkeley, California, United States

Vikash Mansinghka
MIT, Cambridge, Massachusetts, United States

Josh Tenenbaum
MIT, Cambridge, Massachusetts, United States

Armando Solar-Lezama
Massachusetts Institute of Technology, Cambridge, Massachusetts, United States

Abstract

Program induction is an appealing model for human concept learning, but faces scaling challenges in searching the massive space of programs. We propose a computational model capturing two key aspects of human concept learning – our ability to judge how promising a vague, partial hypothesis is, and our ability to gradually refine these vague explanations of observations to precise ones. We represent hypotheses as probabilistic programs with randomness in place of unresolved programmatic structure. To model the evaluation of partial hypotheses, we implement a novel algorithm for efficiently computing the likelihood that a probabilistic program produces the observations. With this, we guide a search process whereby high-entropy, coarse programs are iteratively refined to introduce deterministic structure. Preliminary synthesis results on list manipulation and formal grammar learning tasks show improvements in sample efficiency when leveraging likelihood guidance, and a preliminary human study explores how model intermediate hypotheses compare to those of participants.

In L. K. Samuelson, S. L. Frank, M. Toneva, A. Mackey, & E. Hazeltine (Eds.), *Proceedings of the 46th Annual Conference of the Cognitive Science Society.*