

Modeling Expertise with Neurally-Guided Bayesian Program Induction

Catherine Wong

Massachusetts Institute of Technology, Cambridge, Massachusetts, United States

Kevin Ellis

Massachusetts Institute of Technology, Cambridge, Massachusetts, United States

Mathias Sabl-Meyer

PSL/Collge de France, Paris, France

Josh Tenenbaum

Massachusetts Institute of Technology, Cambridge, Massachusetts, United States

Abstract

Studies of human expertise suggest that experts and novices “see“ problems differently. Experts not only acquire a body of domain-specific strategies and knowledge, but also learn to quickly identify when those concepts apply to problems within the domain. We propose modeling these elements as an iterative process of domain-specific language (DSL) learning, while jointly training a neural network to recognize when learned concepts apply to new problems. We show that the algorithm solves problems more accurately and quickly than either a neural network alone, or a model that simply acquires new concepts without learning when to use them. We also examine the implicit problem representations learned by the neural network recognition model, and find that they increasingly come to reflect abstract relationships between problems, rather than surface features, as the model acquires domain expertise. A full paper and additional details are available at: <https://sites.google.com/view/neurally-guided-expertise-mit>