Decomposed Inductive Procedure Learning: Learning Academic Tasks with Human-Like Data Efficiency

Daniel Weitekamp

[1]Carnegie Mellon University
weitekamp@cmu.edu

Abstract

Human brains have many differently functioning regions which play specialized roles in learning (Poldrack and Foerde 2008). By contrast, methods for training artificial neural networks, such as reinforcement-learning, typically learn exclusively via a single mechanism: gradient descent. This raises the question: might human learners' advantage in learning efficiency over deep-learning be attributed to the interplay between multiple specialized mechanisms of learning? In this work we review a series of simulated learner systems which have been built with the aim of modeling human student's inductive learning as they practice STEM procedural tasks. By comparison to modern deep-learning based methods which train on thousands to millions of examples to acquire passing performance capabilities, these simulated learners match human performance curves—achieving passing levels of performance within about a dozen practice opportunities. We investigate this impressive learning efficiency via an ablation analysis. Beginning with end-to-end reinforcement learning (1-mechanism), we decompose learning systems incrementally to construct the 3-mechanism inductive learning characteristic of prior simulated learners such as Sierra (VanLehn 1990), SimStudent (Matsuda, Cohen, and Koedinger 2015) and the Apprentice Learner Architecture (Maclellan et al. 2016) . Our analysis shows that learning decomposition has a significant effect on the data-efficiency of learning—more so even than simple symbolic/subsymbolic distinctions. Finally we highlight how this breakdown in learning mechanisms can flexibly incorporate diverse forms of natural language and interface grounded instruction (Weitekamp et al. 2023), and discuss opportunities for using these flexible learning capabilities in interactive task learning systems that learn directly from a user's natural instruction (Weitekamp, Harpstead, and Koedinger 2020).

References

Maclellan, C. J.; Harpstead, E.; Patel, R.; and Koedinger, K. R. 2016. The Apprentice Learner Architecture: Closing the Loop between Learning Theory and Educational Data. *International Educational Data Mining Society*.

Matsuda, N.; Cohen, W. W.; and Koedinger, K. R. 2015. Teaching the Teacher: Tutoring SimStudent Leads to More Effective Cognitive Tutor Authoring. *International Journal of Artificial Intelligence in Education*, 25(1): 1–34.

Poldrack, R. A.; and Foerde, K. 2008. Category learning and the memory systems debate. *Neuroscience & Biobehavioral Reviews*, 32(2): 197–205.

VanLehn, K. 1990. *Mind bugs: The origins of procedural misconceptions*. MIT press.

Weitekamp, D.; Harpstead, E.; and Koedinger, K. 2020. An Interaction Design for Machine Teaching to Develop AI Tutors. *CHI*.

Weitekamp, D.; Rachatasumrit, N.; Wei, R.; Harpstead, E.; and Koedinger, K. 2023. Simulating Learning from Language and Examples. In *International Conference on Artificial Intelligence in Education*, 580–586. Springer.