COMPUTER SCIENCE & ENGINEERING
UNIVERSITY OF MICHIGAN

DISSERTATION DEFENSE

Mazin Assanie

Learning General and Correct Procedural Knowledge in a Cognitive Architecture

Tuesday, July 5, 2022
1:00 – 3:00pm
3725 Beyster
Hybrid – Zoom

ABSTRACT: For cognitive architectures that encode procedural knowledge as rules, online procedural learning algorithms exist that can analyze an agent's experiences to acquire additional procedural knowledge. Without careful engineering, agents using existing algorithms can learn large numbers of overly-specific rules that apply in limited situations. Moreover, these algorithms have difficulties summarizing some types of agent reasoning, producing rules that generate unjustified inferences. The limited generality of the rules learned, the uncertainty of correct behavior and the amount of engineering required has limited the effectiveness of online procedural learning (OPL) algorithms. This work provide a detailed analysis of the underlying deficiencies of the most advanced OPL algorithm currently found in a cognitive architecture, Soar's chunking mechanism, by examining two key qualities of OPL: correctness, whether the procedural knowledge generates the same inferences as the reasoning learned from, and optimal generality, whether the procedural knowledge learned captures the generality of the reasoning learned from. This work also provides an architecture-agnostic analytical framework to understand issues with OPL in the form of (1) four necessary conditions for correct behavior summarization, which can be used to organize the correctness issues of an OPL algorithm into a taxonomy, and (2) five strategies that can be used to remedy correctness issues. Finally, this work describes a robust implementation of this approach, a dependency-based procedural learning algorithm called explanation-based behavior summarization (EBBS), which includes mechanisms that remedy or detect each of the correctness issues identified. To improve the generality of procedural knowledge learned, EBBS uses a novel unification algorithm called Distributed Identity Graph Unification (DIGU) that captures the generality in the reasoning being summarized, interfaces with the various mechanisms that improve correctness by using the formalism of object ``identity,'' and is uniquely designed to handle the computational demands that arise from online procedural learning in a cognitive architecture. This work concludes with an evaluation that uses data generated from 14 agents across 8 different domains. These experiments show that EBBS learns more optimally general rules that capture reasoning that previously could not be captured, effects correct behavior and improves agent performance without sacrificing agent reactivity.

CHAIR: Prof. John Laird