Report on the IJCAI-09 Workshop on Learning Structural Knowledge from Observations (STRUCK-09)

Ugur Kuter and Hector Munoz-Avila

■ *The IJCAI-09 Workshop on Learning Structural Knowledge from Observations (STRUCK-09) took place as part of the International Joint Conference on Artificial Intelligence (IJCAI-09) on July 12 in Pasadena, California. The workshop program included paper presentations, discussion sessions about those papers, group discussions about two selected topics, and a joint discussion.*

Human cognition organizes knowledge in different complexity levels: higher-level knowledge is formed by first acquiring simple concepts, which are then combined to learn complex ones. As a result, many cognitive architectures use structural models to represent relations between knowledge of different complexity. Structural modeling has led to a number of representation and reasoning formalisms including frames, schemas, abstractions, hierarchical task networks (HTNs), and goal graphs among others. These formalisms have in common the use of certain kinds of constructs (for example, objects, goals, skills, and tasks) that represent knowledge of varying degrees of complexity and that are connected through structural relations.

In recent years, we have observed increasing interest toward the problem of learning such structural knowledge from observations. These observations range from traces generated by an automated planner to video feeds from a robot performing some actions.

The goal of the workshop was to bring researchers together from machine learning, automated planning, case-based reasoning, cognitive science, and other communities that are looking into instances of this problem and to share ideas and perspectives in a common forum. A total of 27 researchers from

these communities participated in the workshop.

The workshop program included presentation of 12 papers. These were grouped together in four sessions based on the common themes of the workshop: (1) learning structural knowledge from demonstrations (3 papers); (2) learning structure and constraints (4 papers); (3) learning hierarchical planning knowledge (3 papers); and (4) learning structure under uncertainty (2 papers).

Following each session we had a panel discussion about each of these topics in which paper authors and the following session chairs led the discussion: Susan Epstein (CUNY), David J. Stracuzzi (Arizona State University), Santi Ontañón Villar (Georgia Institute of Technology), and Fusun Yaman (BBN Technologies).

Based on the submissions, we identified two overarching themes that were debated in two separate discussion groups:
learning structure under uncertainty and partial information; and structural knowledge and the languages to represent them: trade-offs, capabilities, and limitations. These groups were led by Prasad Tadepalli (Oregon State University) and Daniel Shapiro (ISLE), respectively. In the first discussion group, the problem of how to learn structural knowledge in continuous domains (that is, planning domains that involve numeric variables) emerged as the prominent topic. Questions such as how to represent uncertainty in such domains and how to learn over such representations were discussed in several different planning scenarios. In the second discussion group, the prominent topics discussed included qualitative and quantative representations for learning and trade-offs such as learnability versus expressivity. These discussions led to the question of how much current representations enable human understanding of both the knowledge that is input to the learning system as well as the knowledge that is output by that system. After the group discussions, each of the group coordinators presented a summary of the group discussions.

We are thankful to each of the individuals mentioned before for making the STRUCK-09 workshop a success and to the IJCAI-09 workshop organizers for their support.

Ugur Kuter is an assistant research scientist at the Institute for Advanced Computer Studies, University of Maryland College Park.

Hector Munoz-Avila is an associate professor in the Computer Science and Engineering department at Lehigh University.

Call for Papers:
Sixth
Annual AAAI Artificial
Intelligence for Interactive
Digital Entertainment Conference
(AIIDE-10)

The Sixth Annual AAAI Artificial Intelligence for Interactive Digital Entertainment Conference will be held October 11–13, 2010, at Stanford University adjacent to Palo Alto, California. AIIDE-10 organizers are currently accepting submissions for the Sixth Annual Artificial Intelligence and Interactive Digital Entertainment Conference, intended to be the definitive point of interaction between entertainment software developers interested in AI and academic and industrial AI researchers. While traditionally emphasizing commercial computer and video games, researchers and developers are invited to share their insights and cutting-edge results on all topics at the interface of entertainment and artificial intelligence, including serious games, entertainment robotics, and beyond. Because AIIDE-10 crosses disciplinary boundaries, submissions will be evaluated based on their accessibility to both commercial game developers and researchers in addition to their technical merit.

AIIDE-10 will consist of two tracks, including the Research Track, where papers describe AI research results that make advances towards solving known game AI problems or enabling a new form of interactive digital entertainment, and the Industry Track, developed for individuals that have game development experience but lack the time or need for publishing rigorous academic papers. This track will include presentations of AI techniques, issues, or case studies from the perspective of implementing a product in the current commercial environment.

All submissions are due April 20, 2010. The complete call for papers and other information can be found at the AIIDE-10 web sites: www.aaai.org/aiide10 or www.aiide.org.