

The First Competition on Knowledge Engineering for Planning and Scheduling

Roman Barták and Lee McCluskey

■ We report on the staging of the first competition on knowledge engineering for AI planning and scheduling systems, held in Monterey, California, in colocation with the ICAPS 2005 conference. The background and motivation is discussed, together with the relationship of this new competition with the current international planning competition. We report on the new competition's format, its outcome, and the benefits we hope it will bring to the research area.

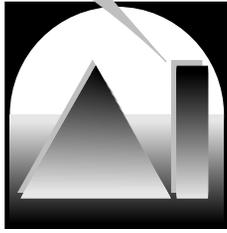
A major goal of AI competitions is to accelerate development in some specified area. Since its first run in 1998, the AI planning community has organized the biannual International Planning Competition (IPC), where various planning systems battle in the field of solving planning tasks. There is no doubt that this competition has brought many benefits to the community. It focused researchers on technology innova-

tion, led to a rapid development of techniques, established a common language for communicating the dynamics of domain models (PDDL), and helped in the validation and scrutiny of planning algorithms. Further, the IPC has facilitated the sharing of benchmark domain models, tasks, and planning tools through the use of PDDL. However, the narrow focus and limiting assumptions of the IPC are controversial—it encourages rapid development, but in the narrow area of fully autonomous plan generation. Competitors taking part in the IPC act on the assumption that the input to their planning engine (a domain model and a planning task) constitutes a bug-free statement of the dynamics of a world that the model represents. The domain models and planning tasks have to be highly refined to ensure that solutions are logically possible. Further, they must be written to satisfy the limiting environmental assumptions of the class of planners that they are testing.

In reality, autonomous plan generation is only one part of systems that

embody planning and scheduling (P&S). The power and efficiency of these systems are also limited by the quality of the application knowledge that they use. If the domain model is flawed, the resulting P&S application will be flawed. And this is the area where the new international competition on knowledge engineering in planning and scheduling (ICKEPS) aims to contribute, in a way complementary to the IPC. Knowledge engineering of planning systems is the process that deals with the acquisition, validation, and maintenance of planning domain models and the selection and optimization of appropriate planning machinery to work on them. Knowledge engineering processes support the planning process: they comprise all of the offline, knowledge-based aspects of planning that are to do with the application being built and any online processes that cause changes in the planner's domain model. The aim of ICKEPS is to promote the knowledge-based aspects of planning by evaluating knowledge engineering tools within a competitive forum. ICKEPS is aimed particularly at the knowledge-based and domain modeling aspects of P&S to accelerate knowledge engineering research in AI P&S and to encourage the development and sharing of prototype tools or software platforms that promise more rapid, accessible, and effective ways to construct reliable and efficient P&S systems. We expect the competition to encourage the development of tools across the whole KE area including domain modeling, heuristic acquisition, planner-domain matching, domain knowledge validation, and so forth.

Within the established IPC, algorithms are evaluated using agreed inputs (domain models, task instances) and standard outputs (plans). Given the limiting assumptions previously outlined, planner evaluation and results production of the IPC can to a large extent be automated. Measures such as "speed of solution" and "quality of solution" can be formulated objectively. With ICKEPS, the difficulty in pinning down inputs and evaluating outputs made the event difficult to stage. Tools and methods to support



*Please Join Us for
IAAI-06*

For details, see
www.aaai.org/Conferences/IAAI/iaai06.php

knowledge acquisition and modeling do not have standard forms of input. They may acquire knowledge from domain experts, or they may help planning researchers debug domain models. Assuming that knowledge engineering tools help create domain models, they cannot be easily evaluated by their outputs—what is the advantage of one domain model over another? Also, tool support for knowledge engineering is heterogeneous: there may be several types of tools performing different functions.

Taking these issues into account, it was decided to stage the first competition with a very simple format and to assemble a panel of judges to evaluate the competing systems. Any system could be entered that fell within the scope of the competition and that was not associated with any of the judges. The scope included tools that helped in knowledge formulation (the acquisition and encoding of domain structure or control heuristics), planner configuration (fusing application knowledge with a P&S engine), validation of the domain model (for example, using visualization, analysis, reformulation) or validation and maintenance of the P&S system as a whole (for example, using plan/schedule visualization, or automated knowledge refinement). The openness of the

competition made the objective evaluation of the tools quite complicated. Instead, through a process of consultation with the organizing committee and interested researchers, we formulated some basic criteria that the judges could use to select the winner. This included support potential, scope of the system, usability, interoperability, innovation, wider comparison with knowledge engineering tools in general AI, build quality, and relevance to the scope.

The actual competition was colocated with ICAPS 2005 in Monterey and ran in two stages. In the preconference stage, the competitors submitted short papers describing the tools. The program committee did light reviewing of the papers with the goal to evaluate relevance of the tools, to send feedback to the competitors, and to contribute to the overall evaluation. During the conference, the competitors gave talks about their systems in a workshoplike arrangement, and then they presented the systems during the ICAPS open demonstration session. The final decision about the winners was done by three judges: Amedeo Cesta (ISTC-CNR, Italy), Sylvie Thiébaux (The Australian National University, Australia), and David E. Wilkins (SRI International, USA).

And what was the result? Seven systems, from around the world, participated in the first run of ICKEPS. These systems contained tools covering a large area of the scope of the competition, in particular knowledge acquisition, knowledge formulation, knowledge modeling, machine learning, knowledge refinement, visualization, analysis, and debugging. In the event the judges split the systems into two categories, and each category had its own winner. Three systems were in the “general tools” category: ModPlan (Germany), GIPO (UK), and ItSIMPLE (Brazil). GIPO by Ron M. Simpson won the award in this category. The second category was “specific tools” and comprised four systems, namely Hamlet (Spain), ARMS (Hong Kong), Tailor (USA), and PlanWorks (USA). ARMS by Kangheng Wu, Qiang Yang, and Yunfei Jiang won the award in the second category. The details of these systems can be found on the competi-

tion website, which is scom.hud.ac.uk/scomt/m/competition/.

The first KE competition raised a high interest in the planning community, and the next competition is planned to be held during ICAPS 2007. The idea is that the knowledge engineering competition will complement the current planning competition and could be held every “odd” year at ICAPS. We hope that future KE competitions will lead to some form of shared communication languages or media for knowledge-based domain models and that they will encourage the development of tools across the whole KE area, as IPC did in the area of planning algorithms.



Roman Barták is an associate professor of computer science at Charles University in Prague, Czech Republic. He obtained a Ph.D degree in computer science from Charles University in 1997. From 1999–2004

he led research activities of Visopt BV where he was a main architect of the scheduling engine developed by this company. He is an author of on-line guides to Prolog programming and constraint programming and is currently teaching courses on automata and grammars, constraint programming, and planning and scheduling. His research focuses on techniques of constraint satisfaction and their applications in planning and scheduling. His e-mail address is roman.bartak@mff.cuni.cz.



Lee McCluskey is a professor of software technology at the University of Huddersfield in the UK. He obtained his PhD degree in the area of machine learning and planning from The City University, London, in 1989.

His research interests include software and knowledge engineering, planning and machine learning. In the 1990s he led a series of projects funded by the UK’s civil aviation which succeeded in using many-sorted first order logic to capture air traffic control separation knowledge for part of the North Atlantic airspace. Lately he has been researching in knowledge engineering for planning, and he is a conference cochair for ICAPS 2006 which is to be held in Cumbria, UK. His e-mail address is t.l.mclcluskey@hud.ac.uk/.