The Knowledge Engineering Review, Vol. 28:2, 117–119. © Cambridge University Press, 2013
doi:10.1017/S0269888912000379
First published online 22 February 2013

Preface to special issue on Knowledge Engineering for Planning and Scheduling

ROMAN BARTÁK[1], SIMONE FRATINI[2] and THOMAS LEO MCCLUSKEY[3]

[1]*Faculty of Mathematics and Physics, Department of Theoretical Computer Science and Mathematical Logic, Charles University in Prague, Malostranské nám. 25, 118 00 Praha 1, Czech Republic;*
e-mail: bartak@ktiml.mff.cuni.cz;
[2]*Institute of Cognitive Science and Technology of the Italian National Research Council, Via S. Martino della Battaglia 44, I-00185 Rome, Italy;*
e-mail: simone.fratini@istc.cnr.it;
[3]*Department of Informatics, University of Huddersfield, Huddersfield, West Yorkshire, HD1 3DH, UK;*
e-mail: t.l.mccluskey@hud.ac.uk

The main focus of Planning and Scheduling (P&S) research is centred around the construction of solver engines, which accept a domain (and task) model as input, and output solutions to P&S problems. This research is beginning to lead to the widespread use of the technology, and P&S solvers can be found embedded in a number of application areas, including control of UAVs, story narrative generation, Web service composition, and logistics. The domain models for these applications are created to contain knowledge such as the physics of the actions, the objects affected by the actions, and the goals or tasks that require solving. The focus on the solver engine needs to be complemented with research on the construction, validation, and optimisation of domain models and domain model languages, such as PDDL (Planning Domain Definition Language). Interest in knowledge engineering (KE) for P&S therefore has grown in recent years, and is concentrated on the formulation of application knowledge bases and the fusion of them with planning engines to create operational systems. Although the main aspect of KE is in domain modelling, it also includes the areas of heuristic acquisition, planner-domain matching, domain knowledge validation, and so forth.

We can categorise methods of producing domain models containing representations of actions, ready for use in a solver, into two main areas:

1. Handcrafted method: a user builds up a model from scratch, perhaps aided with the use of knowledge acquisition tools. In this case, the user needs to be an expert in the domain model language that the solver accepts.
2. Automated method: a user creates or obtains a model in an application-oriented language, and a translator or learning algorithm inputs the model in the application language and outputs solver-ready domain models. The input model could range from being in the form of a declarative formal specification, to a set of example plan scripts.

Over the last 40 years in the development of P&S, domain models have been produced in the main using method 1. An important aspect, relevant for the progression of the field of domain independent P&S, is that general solver engines can be accessed and used by non-AI (artificial intelligence) experts. The development of tools and techniques within method 2 furthers this aim, and overcomes the need to handcraft domain models, which is seen by some as a limiting factor in the deployment of P&S. This recognises that experts in that area may be familiar with their own description languages, but not with P&S description languages such as PDDL. It holds the promise that P&S solvers could be embedded into tool support in the application without the need for a planning expert.

This special issue contains the selected and extended papers of some of the competitors from the third competition on KE for AI P&S systems (International Competition on Knowledge Engineering for Planning and Scheduling (ICKEPS)), held during ICAPS, the International

Conference on Automated Planning and Scheduling at Thessaloniki, Greece in September 2009. ICKEPS promotes the knowledge-intensive aspects of P&S by evaluating KE tools within a competitive forum. The focus of the third competition, and the papers contained in this volume, is on the particular aspect of KE described in 2 above. We believe that this promotes the visibility, usability, and exploitation of P&S solver engines, and could lead to further development of the technology as their use in new applications uncovers new directions and challenges. Papers in this volume deal with P&S application areas such as Web Service Composition, Data Mining, Business Modeling, and E-Learning.

- Antonio Garrido, Susana Fernández, Lluvia Morales, Eva Onaindía, Daniel Borrajo, and Luis Castillo report on the application of AI planning to the area of e-Learning (Garrido *et al.*, 2013). Their system translates course definitions and student information written in an e-Learning language, into solver-ready domain models and planning tasks, respectively. The output plan is translated to a learning design appropriate to the student. While showing the feasibility of the approach, the research also highlights interesting limitations in both e-Learning languages and planning technology.
- The Porsce II framework of Ourania Hatzi, Dimitris Vrakas, Nick Bassiliades, Dimosthenis Anagnostopoulos, and Ioannis Vlahavas is aimed at supporting the process of combining atomic Web services in order to solve a complex goal (Hatzi *et al.*, 2013). Within this context, the system translates models of Web services, written in the Web service ontology language OWL-S, into PDDL, in order that a planner may synthesise a plan that translates back as a composite Web service. The paper details various other functions of Porsce II, which enhance the translations using a mixed-initiative approach.
- In the area of Data Mining, it is well known that the use of mining techniques involves having to extensively pre-process data, set system parameters, and carry out training actions repeatedly over several steps. The system by Susana Fernández, Tomás de la Rosa, Fernando Fernández, Rubén Suárez, Javier Ortiz, Daniel Borrajo, and David Manzano translates the potential steps from a well-known data mining platform into PDDL actions, thus enabling a planner to be used to create plans representing a path in the space of data transformations and training actions (Fernández *et al.*, 2013). Along the way the authors introduce PMML2PDDL, a tool for translating the data mining models written in PMML (Predictive Model Markup Language) to PDDL.
- The contribution of Arturo González-Ferrer, Juan Fernández-Olivares, and Luis Castillo lies in developing a translator to create planning domain models from Business Process Management Notation (BPMN; González-Ferrer *et al.*, 2013). Their system provides a tool for analysts that need to perform resource allocation analysis on business workflows. It embeds a complex transformation of BPMN-expressed workflows in terms of HTNs allowing the exploitation of the ubiquitous BPMN standard for workflow specification.
- The next paper details a system which is rather different from the systems above in that it is application neutral. Learning Object-Centred Models (LOCM) by Stephen N. Cresswell, Thomas Leo McCluskey, and Margaret Mary West inputs sets of example planning scripts, within a well-defined object-centred format, and builds PDDL domain models (Cresswell *et al.*, 2013). At its heart, LOCM is an inductive learning system that induces all the components of a PDDL domain model—predicates, states, and action representations. Rather than relying on knowledge implicit in an input application model, it makes assumptions about the kind of domain model it is building on which to base its generalisations.
- itSIMPLE 3.0 by Tiago S. Vaquero, José R. Silva, Flavio Tonidandel, and J. Christopher Beck is a mature system capable of supporting analysts in modelling domains, planning with them, and visualising the outcomes of planning activities (Vaquero *et al.*, 2013). The system contains user-friendly GUI interfaces that input standard Unified Modeling Language (UML) as the input representation language. Like LOCM, itSIMPLE is not application specific, and the tool has been used in several applications. The choice of UML as input language brings to the tool the potential for widespread use, especially by mainstream software engineers and analysts.

All in all, this volume contains details of exciting research in opening up applications to the use of embedded P&S technology. The contributing authors demonstrated the real bridges between the advanced P&S technology and real-world applications. The guest editors would like to thank all authors for their interesting submissions and also to numerous reviewers who put their effort and time to help the authors with improving their contributions. Special thanks go to the editors-in-chief of KER journal, Simon Parsons and Peter McBurney, for their support with this special issue.

References

Cresswell, S. N., McCluskey, T. L. & Mary West, M. 2013. Acquiring planning domain models using LOCM. *The Knowledge Engineering Review* **28**(2), 195–213.

Fernández, S., de la Rosa, T., Fernández, F., Suárez, R., Ortiz, J., Borrajo, D. & Manzano, D. 2013. Using automated planning for improving data mining processes. *The Knowledge Engineering Review* **28**(2), 157–173.

Garrido, A., Fernández, S., Morales, L., Onaindía, E., Borrajo, D. & Castillo, L. 2013. On the automatic compilation of e-learning models to planning. *The Knowledge Engineering Review* **28**(2), 121–136.

González-Ferrer, A., Fernández-Olivares, J. & Castillo, L. 2013. From business process models to hierarchical task network planning domains. *The Knowledge Engineering Review* **28**(2), 175–193.

Hatzi, O., Vrakas, D., Bassiliades, N., Anagnostopoulos, D. & Vlahavas, I. 2013. The PORSCE II framework: using AI planning for automated Semantic Web service composition. *The Knowledge Engineering Review* **28**(2), 137–156.

Vaquero, T. S., Silva, J. R., Tonidandel, F. & Beck, J. C. 2013. itSIMPLE: towards an integrated design system for real planning applications. *The Knowledge Engineering Review* **28**(2), 215–230.