

AAAI-10 Classic Paper Award

Systematic Nonlinear Planning: A Commentary

Daniel S. Weld

■ David McAllester and David Rosenblitt's paper, "Systematic Nonlinear Planning" (published in the Proceedings of the Ninth National Conference on Artificial Intelligence [AAAI-91]), won the AAAI-10 classic paper award. This commentary by Daniel S. Weld describes the two major impacts the paper had on the field of automated planning.

The field of automated planning seeks to develop fast methods for generating programs of action that allow an agent to achieve goals or maximize rewards. Initially, researchers made many simplifying assumptions, defining the *classical* planning problem: produce a sequence of atomic actions that will achieve a logically specified goal in a completely known world where action effects are certain. David McAllester and David Rosenblitt's paper, "Systematic Nonlinear Planning" (McAllester and Rosenblitt 1991), presented 19 years ago at the Ninth National Conference on Artificial Intelligence (AAAI-91), had two major impacts on the field: (1) an elegant algorithm and (2) endorsement of the *lifting* technique. The paper's biggest impact stems from its extremely clear and simple presentation of a sound and complete algorithm (known as SNLP or POP) for classical planning. While it is easy to define such an algorithm as search through the space of world states, SNLP is a "partial-order" planner, meaning it searches the space of partially specified plans, where only partial constraints on action arguments and ordering decisions are maintained. Here, McAllester and Rosenblitt benefited from David Chapman's elegant TWEAK planner, which greatly clarified previous partial-order algorithms (Chapman 1985). SNLP's key feature is the use of a data structure, called a *causal link*, to record the planner's commitment to establish a precondition of one action with the postcondition of another. While the notion of causal link wasn't novel, originating much earlier in Austin Tate's NONLIN planner (Tate 1977), SNLP was so much simpler than NONLIN that McAllester and Rosenblitt could prove it was sound and complete. A bonanza of research results followed. Penberthy and Weld's (1992) UCPOP planner handled actions modeled in an

expressive language and was widely distributed. Other extensions included temporal planning, planning with metric resources, contingent planning in response to uncertain effects, methods for guiding the planner's nondeterministic search, and so on.

A second contribution of McAllester and Rosenblitt's paper was its clear realization that partial-order planning was *not* aimed at finding partially ordered plans, but rather *lifting* the underlying search from operator sequences to *equivalence classes* over operator sequences (represented by a set of ordering constraints). This insight allowed the authors to steer away from "necessary and sufficient conditions for correctness of partially ordered plans," as formalized by Chapman as the *modal truth criterion*, and to provide conditions under which the search is *systematic*, that is, that the equivalence classes don't overlap.

The desire to speed planning led many to abandon partial-order planning in the late 1990s. The impetus for change came from three alternative approaches to planning, each of which appeared to dominate the SNLP family. Avrim Blum and Merrick Furst's Graphplan (Blum and Furst 1995) was the first on the scene. Graphplan operates in two phases — first, in polynomial time it constructs a leveled graph whose nodes are either actions or propositions and whose edges are akin to causal links. In the second phase, Graphplan searches backwards from the graph's final level, looking for a consistent set of actions to achieve the goal. The following year, two more planning algorithms diverted attention from partial-order planners. Using a Graphplanlike, leveled-graph model, SATPLAN (Kautz and Selman 1996) showed that planning problems could be compiled into SAT problems, which modern SAT solvers could solve to create plans, blazingly quickly. Finally, Drew McDermott's UNPOP (McDermott 1996) introduced a method for automatically generating domain-independent heuristics, and extensions of the method allowed simple state-space planners to solve larger problems than previously believed possible. Taken together, these exciting new directions led many researchers away from the SNLP lineage of planners.

But were researchers correct to abandon partial-order planning? While the idea of representing a plan as a partially specified and partially ordered set of actions embodied the popular ideas of least commitment and constraint posting, it had become apparent that the representation harbored an inherent weakness — there was no fast and easy way to compute the world state at a "point" in the plan. Heuristically estimating the distance, that is, the number of plan modification operations from one partial plan to another, that solved the goal, was extremely tricky. In contrast, if one represents

a plan with a totally ordered sequence of concrete actions (as with McDermott's UNPOP), then it is easy to calculate the resulting state and thus heuristically estimate the distance to a desired goal state. In time researchers would recognize that a major component of Graphplan's speed could be attributed to the planning graph's role as a fast way to estimate *reachability* in the state space. This realization led to even faster planners that used a planning graph to quickly generate heuristic values for use in regression and progression planning.

Eventually, researchers took the new understanding of automated heuristic construction back to the space of partially ordered plans, completing the cycle. XuanLong Nguyen and Subbarao Kambhampati's RePOP (Nguyen and Kambhampati 2001) sparked renewed interest in SNLP-style planning by introducing domain-independent heuristics that rendered an SNLP-like algorithm competitive with the state of the art. Interestingly, some of these heuristics were computed from a Graphplan-style planning graph, while others were derived from graph cuts through the partial plan's causal link structure, affording the equivalent to a view of intermediate states during execution of a potential linearization of the partially ordered steps. Then, in the 2002 International Planning Competition, two planners using partial-order representations surprised the community with their performance. Indeed, one of them, Gerevini and Serina's LPG (Gerevini and Serina 2002) won the competition.

Up until now discussion has been restricted to classical planning, but partial-order planning has many advantages in two currently active research areas: temporal planning and integrated planning and execution. Temporal planning abandons the classical assumption that actions execute atomically, instead modeling the expected time duration required for execution. In this framework a world-state representation creates many more irrelevant distinctions (for example, depending on the order in which overlapping actions are started); in contrast, the causal link representation extends naturally. Indeed, most work on temporal planning originally used this representation. While, recent temporal competitions have been dominated by so-called decision epoch planners, work by William Cushing, Daniel Weld, Subbarao Kambhampati, Mausam, and Kartik Talamaduplula (Cushing et al. 2007) shows that they suffer from fundamental problems of expressiveness. Furthermore, when a planner is connected to physical effectors and must recover from execution errors and unexpected outcomes, a partially ordered, causal-link representation allows more flexibility to the executive, and allows for quicker recovery if replanning is required. These reasons likely explain why NASA researchers chose a partial-order representation for EUROPA, which planned actions for

the Mars rovers *Spirit* and *Opportunity* (Bresina et al. 2005).

References

Blum, A., and Furst, M. L. 1995. Fast Planning through Planning Graph Analysis. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence, IJCAI 95*, 1636–1642. San Francisco: Morgan Kaufmann Publishers.

Bresina, J. L.; Jónsson, A. K.; Morris, P. H.; Rajan, K. 2005. Activity Planning for the Mars Exploration Rovers. In *Proceedings of the Fifteenth International Conference on Automated Planning and Scheduling (ICAPS 2005)*, 40–49. Menlo Park, CA: AAAI Press.

Chapman, D. 1985. Nonlinear Planning: A Rigorous Reconstruction. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence (IJCAI-85)*, 1022–1024. San Francisco: Morgan Kaufmann Publishers.

Cushing, W.; Weld, D. S.; Kambhampati, S.; Mausam; Talamadupula, K. 2007. Evaluating Temporal Planning Domains. In *Proceedings of the Seventeenth International Conference on Automated Planning and Scheduling, (ICAPS 2007)*, 105–112. Menlo Park, CA: AAAI Press.

Gerevini, A., and Serina, I. 2002. LPG: A Planner Based on Local Search for Planning Graphs with Action Costs. In *Proceedings of the Sixth International Conference on Artificial Intelligence Planning Systems (AIPS 2002)*, 13–22. Menlo Park, CA: AAAI Press.

Kautz, H. A., and Selman, B. 1996. Pushing the Envelope: Planning, Propositional Logic, and Stochastic Search. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence and Eighth Innovative Applications of Artificial Intelligence Conference, AAAI 96*, 1194–1201. Menlo Park, CA: AAAI Press.

McAllester, D., and Rosenblitt, D. 1991. Systematic Non-linear Planning. In *Proceedings of the Ninth National Conference on Artificial Intelligence (AAAI-91)*. Menlo Park, CA: AAAI Press.

McDermott, D. V. 1996. A Heuristic Estimator for Means-Ends Analysis in Planning. In *Proceedings of the Third International Conference on Artificial Intelligence Planning Systems*, 142–149. Menlo Park, CA: AAAI Press.

Nguyen, X., and Kambhampati, S. 2001. Reviving Partial Order Planning. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence, (IJCAI 2001)*, 459–466. San Francisco: Morgan Kaufmann Publishers.

Penberthy, J. S., and Weld, D. S. 1992. UCPOP: A Sound, Complete, Partial Order Planner for ADL. In *Proceedings of the Third International Conference on Principles of Knowledge Representation and Reasoning (KR'92)*, 103–114. San Francisco: Morgan Kaufmann Publishers.

Tate, A. 1977. Generating Project Networks. In *Proceedings of the Fifth International Joint Conference on Artificial Intelligence (IJCAI-85)*, 1022–1024. Los Altos, CA: William Kaufmann Inc.

Daniel S. Weld is the Thomas J. Cable/WRF Professor of Computer Science and Engineering at the University of Washington, Seattle, Washington, USA.

CAMBRIDGE


NOTEWORTHY TITLES from CAMBRIDGE!

Logical Dynamics of Information and Interaction

JOHAN VAN BENTHEM

Presents logic as a way to identify all information processes that drive human action.


\$95.00; Hardback: 978-0-521-76579-4; 384 pp.



Machine Ethics

Edited by MICHAEL ANDERSON
and SUSAN LEIGH ANDERSON

\$99.00; Hardback: 978-0-521-11235-2; 480 pp.




Face Geometry and Appearance Modeling Concepts and Applications

ZICHENG LIU

and ZHENGYOU ZHANG

\$80.00; Hardback: 978-0-521-89841-6; 320 pp.




Graph-based Natural Language Processing and Information Retrieval

RADA F. MIHALCEA

and DRAGOMIR R. RADEV

This book extensively covers the use of graph-based algorithms for natural language processing and information retrieval. Readers will come away with a firm understanding of the major methods and applications of these topics that rely on graph-based representations and algorithms.

\$65.00; Hardback: 978-0-521-89613-9; 224 pp.



Research Methods for Science

MICHAEL P. MARDER

Unique introduction to the design, analysis, and presentation of scientific projects, an essential textbook for undergraduates in science and mathematics.

\$25.99; Paperback: 978-0-521-14584-8; 236 pp.



Prices subject to change.



CAMBRIDGE
UNIVERSITY PRESS

www.cambridge.org/computerscience
800.872.7423