Don't Leave Your Plan on the Shelf

Austin Tate
Artificial Intelligence Applications Institute
The University of Edinburgh, 80 South Bridge, Edinburgh EH1 1HN, UK

From Symposium on Artificial Intelligence, Special Issue of *ACM Computing Surveys*, (eds. J.Doyle and P.Wegner), Association of Computing Machinery, September 1995.

The field of AI planning started off well. Over 25 years ago there was an AI planning system based on well-founded concepts that could generate plans for an automated robot and modify those plans in the face of execution uncertainty and failure (in the work of Fikes, Nilsson and their colleagues at Stanford and SRI). Researchers of the time were exploring the common requirements for automated programming, natural language understanding, constraint handling, human operator guidance and robot planning (in the work of Winograd, Sussman and Winston at MIT and Waldinger and Sacerdoti at SRI). We are only now returning to the same level of aspiration as these early researchers. What happened in between?

1. Planning in general is hard.
2. Toy problems and puzzles were the wrong things to work on.
3. The power of using knowledge about a domain went unrecognized.
4. Study of search issues and the formal properties of search spaces dominated.
5. The context within which planners operated was ignored.

I expand on these points below.

First, planning in the general case is a computationally complex problem. In the early 1970s there was a growing realization that some fundamental problems of plan representation and reasoning (abstraction, hierarchies of task networks, causal structure, resources, time constraints, etc.) had to be addressed for progress to be made. Each of these was to involve many research efforts. This tended to fragment overall visions of realistic planning systems.

Second, a number of puzzles had been identified by the early researchers as problematic for their systems, methods or representations - such as the three blocks problem or the keys and boxes problem. Unfortunately, rather than being used for clear exposition of new methods these became a topic for study in their own right. The formal properties of these puzzles lend themselves to thorough analysis, but unfortunately, they are far from perfect substitutes for the real problems that need addressing for AI planning techniques to be useful.

Third, the lessons of the late 1970s in making good use of knowledge about a domain for tasks such as analysis, interpretation and diagnosis (in the so called "expert systems") were late coming to the AI planning community. Only a few years ago, you could still hear cries of "cheat" or "you are building in the solution" if rich models of the problem domain were used in planning systems.

Fourth, interest (and publications) in the field for much of the 1980s and early 1990s was dominated by those concerned with formal characterization of the search spaces of systems only able to deal with the simplest puzzles. These same systems, in seeking a formal basis, did not attempt to model the more "esoteric" features of practical planners like hierarchical task network expansions, domain-knowledge constrained options, rich resource models, temporal information, environmental context, etc. If a real problem was set up in the general frameworks proposed, the search spaces were unrealistic.

Finally, planners were studied and developed in isolation and had a simple notion of the way in which they were tasked and the results used. Work in the late 1980s and onwards started to explore a very different style of situated planning system for reactive plan execution, but it took some time to merge this work with the separate generative planning concepts. There is still little work in the AI planning community on the command or tasking interface to planners - yet it is from this level that the constraints on what can be generated are identified and where result quality issues need to be negotiated.

But the dark ages are over. There is a broader understanding of the formal properties of planners and plans. There is a growing convergence of theoretical work in AI planning and the modeling used in practical planning systems (Weld, Yang, Tate). Some comparative models are now available that allow a broad range of planning systems to be characterized and studied (McDermott, Hendler, Khambhampati). There are effective links between pro-active planning systems and reactive execution support systems. A start is being made on learning some of the lessons from knowledge engineering and acquisition and applying them to the field of plan and activity management (Valente, Gil, Chien). Even more encouragingly, there is a growing (renewed) convergence in the techniques used in the planning field with those used in automatic programming, process management, computer-aided software engineering, operations research, business modeling, cooperative working and workflow support.

The common ground for the future may be a better formulation of and understanding of the nature of plans themselves. This would include their initial and changing requirements and would reflect the environment within which they are being, or are intended to be, executed. AI planning has much to offer other areas of information technology. The importance of knowledge-rich plan representations for work in systems design, process management, cooperative working and workflow systems has still to be realized. Let's make sure we don't leave our plans on the shelf.

Reference:

Readings in Planning, Allen, J., Hendler, J. and Tate, A. (editors), Morgan Kaufmann Publishers, Inc., San Mateo, California, 1990.