The Knowledge Engineering Review, Vol. 13:1, 1998, 179–184

Robots, softbots, immobots: The 1997 AAAI Workshop on Theories of Action, Planning and Control

ROBERT P. GOLDMAN[1] and CHITTA BARAL[2]

[1]*Honeywell Technology Center, 3660 Technology Drive Minneapolis, MN 55418. Email: goldman@htc.honeywell.com*
[2]*Department of Computer Science, University of Texas at El Paso, El Paso TX 79968, USA. Email: chitta@cs.urep.edu*

1 Introduction

This workshop brought together researchers concerned with fundamental issues of modelling action, those developing automated planning techniques and those attempting to implement autonomous agents acting in the world. In the past, these research communities have been separate from each other to a surprising extent. Researchers interested in theories of action have busied themselves with finding solutions to the frame and ramification problems, for very expressive theories of action. On the other hand, researchers interested in developing planning systems have typically concentrated on efficiency over expressiveness and assumed away the frame and ramification problems by means of the "STRIPS assumption".[1] Finally, researchers interested in implementing autonomous agents have found their attention occupied by issues of execution monitoring and sensing.

The 1997 AAAI workshop was the successor to a similar workshop held at AAAI-96. At that workshop there was interesting interaction between planning researchers and researchers interested in theories of action. A particularly active topic of discussion concerned planning and action under conditions of incomplete information. At that time there had been widespread interest in both research communities on action under uncertainty.

At the 1996 workshop, although there was contact between planning researchers and those developing theories of action, there was less contact with those interested in implementing automated agency. We tried to bridge this gap in the 1997 workshop. In particular, we wished to bring into the discussion the three "bots" that are currently the focus of much discussion among AI researchers: robots, softbots and immobots.

2 The three bots

AI researchers have always been interested in controlling and planning for mobile robots. Recently, however, interest in automated agency has extended to two additional kinds of "bot": "softbots" and "immobots."

The notion of softbot was introduced by Oren Etzioni (1993), at a time of grave concern about the realism of AI research (particularly in the area of planning). Etzioni argued that automated agents operating in a wholly software environment would provide a worthy testbed for AI theories. Adopting softbots as test platforms would have the advantage of offering readier access to real autonomous action than mobile robots; mobile robots are expensive both in terms of raw cost and

[1]That any proposition not explicitly indicated as being changed by an action will persist through that action.

in terms of the amount of effort required to master their mechanics, etc. Softbots would provide a quicker path to deploying automated agents without the drawbacks of the conventional approach, robot simulators. Unlike a robot simulator, which is an idealization of a physical environment, "Software environments are not idealizations of physical environments; developing softbots is a difficult and exciting challenge in its own right." (Etzioni, 1993, p. 9).

More recently, Williams and Nayak (1996) have proposed "immobots" as a third class of autonomous agents. They use this term to describe complex, autonomous systems whose operations are concerned more with their own internal state than with the issues raised by motion and navigation. Williams and Nayak offer large buildings and deep space probes as examples of immobots.

3 Structure of the workshop

In order to focus our discussion, we organized the day around three invited talks and one panel. Following each invited talk was a series of very short presentations by workshop attendees. These were chosen to address some theme struck by the invited speaker.

The three invited talks were given by researchers working in automated agency. We invited these speakers to do short tutorials. In these, the speakers outlined their domains of interest, and discussed key research challenges, rather than giving narrow accounts of their own work.

The first presentation was by Keith Golden, who discussed softbots' requirements for action representation and planning (Golden & Weld, 1996). The second presentation was given by Kurt Krebsbach, who discussed a particular kind of immobot: the control system for an oil refinery. The final invited talk was by Barney Pell, of NASA Ames Research Center. He provided a number of interesting challenge problems for automated agency arising out of the needs of autonomous spacecraft.

4 Softbots

Keith Golden discussed of the limitations of existing planner representations, particularly where information-gathering goals are concerned. While these issues have arisen with particular urgency in planning for softbots, they are of concern in all fields of projective planning under incomplete information.

For example, the action 'ls' gives the agent performing the action information about the files in the directory where the action is performed. Unlike actions which change the world this action does not change the world, rather it changes the information that agent has about the world.

5 A refinery immobot

The second presentation was given by Kurt Krebsbach, who discussed a particular kind of immobot: the control system for an oil refinery. Krebsbach, with colleagues at Honeywell Technology Center, has been working on a project for managing abnormal situations in refineries. In his talk, Krebsbach identified the key need in his domain as being rapid reaction to situations based on extensive knowledge of established procedures. He argued that existing model-based approaches (including conventional projective planning) were not suitable for such problems, and suggested that reactive languages like PRS (Georgeff & Lansby, 1988) and RAPS (Firby, 1987) are more suitable.

Krebsbach began by giving some background on refineries[2]. One very important feature of refineries is that—designed to operate in a highly competitive, low-margin business—they have been extensively optimized for energy and raw material use. For example, often streams of materials that must react at high temperature are preheated by running them through heat exchangers with high

[2]For more information about refineries, we suggest Leffler (1985).

temperature streams that must be cooled. This kind of optimization makes all aspects of refinery operation intertwined. Such intricate relations make simplifications for troubleshooting, like those proposed by Williams and Nayak (1996), which rely on identifying those systems that are causally "downstream" from a failure, infeasible. Furthermore, the mathematical models developed by refiners for control and planning are very computationally expensive and cannot be run fast enough in times of crisis. Finally, even if these models could be run fast enough, they would not provide useful guidance. These models are typically guides to the steady-state operation of the system, rather than the dynamic behavior of the system[3].

Because of these features of refinery operations, Krebsbach argued for the use of special languages suited to reactive execution, rather than projective planning. Such languages permit designers to capture the procedural knowledge that is used now by humans in reacting to emergencies in high-stakes, complex environments like oil refineries.

Krebsbach discussed some other challenge problems that arise in the context of handling refinery emergencies:

- Interacting goals of maintenance and achievement: Keep the plant safe, then make it profitable;
- Delayed Action Effects. There are often very long latencies between when an action is initiated and when it has its effects. For example, cutting feed to a unit may take 10–20 minutes to complete. This makes it difficult to monitor plans to see if they are effective;
- Crisis management systems are of necessity mixed-initiative systems. How are users to monitor the execution of complex reactive plans? In reactive systems, the lack of projection makes it difficult to tell the user "what the system wants to do next".

6 Space applications

The final invited talk was by Barney Pell, of NASA Ames Research Center. He provided a number of interesting challenge problems for automated agency arising out of the needs of autonomous spacecraft. Autonomous spacecraft are becoming ever more important to NASA, following a strategic shift to carrying out more and cheaper missions (Pell et al., 1998; Pell et al., 1997; Chien et al., 1997). Pell discussed several challenge problems for AI that arise out of the particular characteristics of space exploration. These challenge problems include:

- contingency or robust plans
- contingent execution
- diagnosing faults, then having to replan
- planning to plan
- planning to talk
- multi-spacecraft constellations
- system fault protection.

Contingency handling is an important aspect of planning and execution for autonomous spacecraft, arising both in preplanning and execution of plans. Spacecraft must be able to execute contingent plans to react to eventualities such as system failures. However, reaction is not sufficient, because often it is necessary to *preposition* assets in order that one be prepared to react to contingencies. For example, Erann Gat notes that in order to be able to react appropriately to a guidance system failure in the Cassini mission, it is necessary to warm up backup systems well in advance of the failure (Gat, 1996; Musliner & Goldman, 1997).

Remote, resource-limited autonomous spacecraft are forced to confront problems of meta-planning (Boddy & Dean, 1994). For example, remote spacecraft must plan ahead in order to have the computational resources necessary for planning ("planning to plan"). In future multiple-

[3]Some control models cover dynamic behavior, but often do not extend to behaviors that occur in abnormal situations.

spacecraft and satellite systems, it will be necessary to plan in advance to communicate with ground stations ("planning to talk"). When such systems are assembled into constellations they must also address issues of information sharing, teamwork and other multi-agent problems

Another key concern of spacecraft is maintaining their operational status, and operating successfully in less-than-perfect repair. Autonomous spacecraft must be able to detect and repair faults, to the maximum extent possible. When this fails, as in long missions it must eventually, the spacecraft should be able to reconfigure themselves to function as well as possible, even in a degraded state. If enough goes wrong, the system may need to reboot itself, whereupon the agent must wake up, try to diagnose failures and figure out how to continue the mission.

7 Planning, action models and discrete event control

The panel was a discussion of the relation of AI action representations (including those used by planners) to discrete event control, as developed within control theory (Ramadge & Wonham, 1989). It was chaired by Kutluhan Erol, and the panelists were Chitta Baral, Robert Goldman and Froudald Kabanza. Baral stressed the importance of formulating correctness of reactive control for given achievement and/or maintainance goals with respect to action theories, and automatically constructing individual control rules. Kabanza discussed the importance of using temporal logic to specify the goals of an agent and Goldman focussed on the real-time aspects of reactive control.

While Baral separated the agent's actions from the exogenous actions, Kabanza put them together as if the agent's actions have non-deterministic effects. Baral's focus was in defining "correctness" conditions for both individual control rules as well as the control module as a whole. Goldman briefly discussed the CIRCA architecture and its usefulness in specifying real-time aspects of reactive control.

8 HTN planning

In his presentation, given shortly after Krebsbach's talk, Dana Nau presented a system based on hierarchical task network (HTN) planning. Nau has used HTN planning in the development of systems that plan the manufacture of mechanical-electrical components simultaneously with designing them and in a champion bridge-playing system. Nau also argued against conventional planning approaches but his position was less extreme than Krebsbach's. Nau argues that a problem with conventional planning systems is the reliance on models of how operators affect individual features of the domain. Such models are not readily available in domains like machining of parts, because the operators are very complex. Nau solves such problems using HTNs: the HTN structure provides the top-down search guidance and the planner can limit itself to forward simulation of fleshed-out action sequences.

After Nau's presentation there was a discussion between Nau and Baral on whether HTN descriptions can be shown correct with respect to action theories. While Nau believes that HTN is an independent formalism, Baral believes that HTN programs can be compiled or at least shown correct with respect to action theories, similar to his work on formalizing and showing correctnes of reactive control with respect to action theories.

9 Planning as theorem-proving

The presentation that aroused the most discussion among participants was Murray Shanahan's talk on his current work on planning in the Event Calculus. This talk was substantially a defense of planning as theorem-proving. This was an interesting development, coming as it did just before Wolfgang Bibel's (1997) similar defense in his invited talk at IJCAI-97. Shanahan discussed a particular kind of planning as theorem-proving, based on the Event Calculus.

Shanahan's talk raises the interesting question of what is meant by the claims of superiority of planning as theorem-proving. In his paper, he writes that

> Many of the computational concepts central to the literature on partial order planning, such as threats, protected links, promotions and demotions . . .turn out to have direct counterparts in the theorem proving process. It's interesting to note that these features of the logic programming implementation weren't designed in. . . .So our attempt to provide a mathematically respectable answer to the question 'What is partial order planning?' inadvertently offers similar answers to questions like 'What are protected links?'

The interesting thing about this claim is that those features he cites were designed into the planners he cites (TWEAK (Chapman, 1987) and UCPOP (Penberthy & Weld, 1992)), in order that their proofs of soundness and completeness would go through. So these features are not actually so mysterious—indeed, this claim is, in some sense, a claim for preferring a computation that is a proof over a computation that has a proof (of soundness and completeness)[4].

Bibel argues for planning as theorem-proving on the grounds that it will mesh well with other AI components that are also based on deduction. He writes:

> I am . . .one of those who strongly believe that only through a rather uniform approach to any of these different facets [of reasoning] can we ever hope to accomplish systems that are able to do more than "just" playing chess on a worldmaster level (but nothing else) or prove open problems like Robbins' one as done recently (but again nothing else) etc. It is an illusion to think we could just combine all these niche systems to get out something like a general intelligence. Rather the entire approach must be a more universal one from the very beginning. And if so only through a uniform approach could the enormous complexity be overcome[5].

Another, related, claim that Shanahan advanced in his talk was a defense of general-purpose, expressive action representations over action representations chosen specifically to permit efficient planning. This aroused a lot of discussion, coming hard on the heels of Keith Golden's discussion of tradeoffs made in designing a language for efficient planning under incomplete information! Another interesting juxtaposition offered by the workshop was the talks by Krebsbach and Pell which went further and argued for representations that were procedural rather than declarative—going far beyond restricting the expressivity of a language that was still, fundamentally, declarative.

10 Concluding remarks

One possibly disappointing feature of the workshop was the comparative lack of those interested in probabilistic approaches to the problems at hand. Recently, Markov Decision Processes (and Partially Observable Markov Decision Processes) have become a popular semantic model in the planning community (Dean, 1994). As a semantics for planning, Markov Decision Processes have the advantage of providing a principled treatment of uncertainty and supporting decision theoretic notions of optimality. We reluctantly decided to exclude these approaches from our discussion, because of the difficulties of including yet another set of terms into what was intended to be a general discussion.

Since people working on action theories (with some exceptions, e.g. Shanahan's Event Calculus planner, (Bonet et al., 1997)) have not focused on planning issues there is a still a big gap between the theories of actions and the languages used in implemented planning systems. Further, as Krebsbach and Pell showed in their workshop presentations, there is even a further distinction between languages for execution and for planning. It is our hope that workshops like this one, by fostering contact between these research communities, can help close the gap.

The papers in the workshop are currently available at `http://cs.utep.edu/actions/aaai97/`.

[4]These comments by the first author; Baral somewhat differs. Marc Friedman makes a similar point in the ongoing debate about Wolfgang Bibel's paper in *ETAI News Journal and Newsletter on Reasoning about Actions and Change*; see `http://www.ida.liu.se/ext/etai/actions/njl/97003/`.
[5]Bibel's comments appear in a recent issue of the *ETAI News Journal and Newsletter on Reasoning about Actions and Change*; `http://www.ida.liu.se/ext/etai/actions/indexframe.html`

References

Bibel, W, 1997. "Let's plan it deductively!" Unpublished manuscript based on invited address to IJCAI-97, March.

Boddy, M and Dean, T, 1994. "Decision-theoretic deliberation scheduling for problem solving in time-constrained environments" *Artificial Intelligence*.

Bonet, B, Lerincs, G and Geffner, H, 1997. "Proceedings of the Fourteenth National Conference on Artificial Intelligence. A robust and fast action selection mechanism for planning" 714–719. AAAI Press/MIT Press.

Chapman, D, 1987. "Planning for conjunctive goals" *Artificial Intelligence* **32**(3) 333–378.

Chien, S, DeCoste, D, Doyle, R and Stolorz, P, 1997. "Making an impact: Artificial intelligence at the Jet Propulsion Laboratory" *AI Magazine* **17**(1) 103–122.

Dean, T, 1994. "Decision-theoretic planning and Markov decision processes" Unpublished.

Etzioni, O, 1993. "Intelligence without robots: A reply to Brooks" *AI Magazine* **14**(4) 7–13.

Firby, RJ, 1987. "An investigation in reactive planning in complex domains" *Proceedings AAAI-87* 196–201.

Gat, E, 1996. "News from the trenches: An overview of unmanned spacecraft for AI" In I, Nourbakhsh, (ed.), *AAAI Technical Report SSS-96–04: Planning with Incomplete Information for Robot Problems* American Association for Artificial Intelligence, March 1996. (Available at `http://www-aig.jpl.nasa.gov/home/gat/gp.html`)

Georgeff, M and Lansky, A, 1986. "Procedural knowledge" *IEEE Special Issue on Knowledge Representation* **74** 1383–1398, October.

Golden, K and Weld, D, 1996. "Representing sensing actions: The middle ground revisited" *Proceedings of KR-96*.

Leffler, WL, 1985. *Petroleum Refining for the Nontechnical Person*. PennWell Books, Tulsa, OK.

Musliner, DJ and Goldman, RP, 1997. "CIRCA and the Cassini Saturn orbit insertion: Solving a prepositioning problem" *Working Notes of the NASA Workshop on Planning and Scheduling for Space (to appear)*.

Pell, B, Bernard, DE, Chien, SA, Gat, E, Muscettola, N, Pandurang Nayak, P, Wagner, MD and Williams, BC, 1998. "An autonomous spacecraft agent prototype" Autonomous Robotics (to appear).

Pell, B, Gat, E, Keesing, R, Muscettola, N and Smith, B, 1997. "Robust periodic planning and execution for autonomous spacecraft" *Fifteenth International Joint Conference on Artificial Intelligence*.

Penberthy, JS and Weld, DS, 1992. "UCPOP: a sound, complete, partial order planner for ADL" In B, Nebel, C, Rich and W, Swartout, (eds), *Principles of Knowledge Representation and Reasoning Proceedings of the Third International Conference* 103–114. Los Altos, CA. Morgan Kaufmann.

Ramadge, PJG and Wonham, WM, 1989. "The control of discrete event systems" *Proceedings of the IEEE* **77**(1) 81–98.

Williams, BC and Pandurang Nayak, P, 1996a. "Immobile robots: AI in the new millennium" *AI Magazine* **17**(3).

Williams, BC and Pandurang Nayak, P, 1996b. "A model-based approach to reactive self-configuring systems" *Proceedings of the Thirteenth National Conference on Artificial Intelligence* 971–978. Menlo Park, CA. AAAI Press/MIT Press.