Position Paper: Topics for Future Planning Competitions

Sven Koenig
College of Computing
Georgia Institute of Technology
Atlanta, GA 30332-0280
skoenig@cc.gatech.edu

Abstract

In this position paper, we suggest two possible topics for future planning competitions, namely planning with realistic planning objectives and re-planning for sequences of similar planning problems. Both topics appear to be important but currently neglected in the ICAPS community. We feel that the planning competition could set worthwhile research agendas in artificial intelligence planning by adopting these topics.

Introduction

There are several ways of choosing topics for planning competitions. In the following, we discuss two of them:

- First, planning competitions could form around existing research programs that are mature enough to warrant competitions to be able to compare existing planners. In this case, there are already research programs that fit the topic of the competitions and it is thus easy to solicit a sufficient number of entries for them. The competitions then focus on developing better planners for established classes of planning problems, for example, planners with a smaller run time or better plan quality.
- Second, planning competitions could set completely new research agendas, for example, focus on more realistic planning problems than those typically studied in the ICAPS community. In this case, there are no research programs yet that fit the topic of the competitions. Thus, it is easy for newcomers to enter them but their organizers need to determine worthwhile classes of planning problems.

So far, the planning competitions have successfully walked a middle ground between the two extremes: they have started with existing research programs but eventually augmented the planning problems to be solved. For example, they started with deterministic STRIPS planning but eventually augmented it to include real-valued resources. In the following, we suggest two possible topics for future planning competitions, namely planning with realistic planning objectives and re-planning for sequences of similar planning problems. Both topics appear to be important but currently neglected in the ICAPS community. We feel that the planning competition could set worthwhile research agendas in artificial intelligence planning by adopting these topics (although we do not intend to enter the planning competition in 2004 ourselves).

More Realistic Planning Objectives

In deterministic domains, planners from artificial intelligence have traditionally been used with the objective to find any plan that achieves the goal. To make their planning objectives richer, planners then began to associate execution costs with plans and preferred plans that achieve the goal with minimal execution cost, that is, minimal consumption of a limited resource such as time, energy, or money. In probabilistic domains, planners usually either minimize the average execution cost or, if the goal cannot be achieved for sure, maximize the probability of goal achievement. However, these planning objectives are often too simplistic to model the preferences of human decision makers adequately. For example, human decision makers trade-off between the consumption of different resources (Drabble, Koehler, & Refanidis 2002), need to deal with hard and soft deadlines (Dean, Firby, & Miller 1988; Haddawy & Hanks 1992), and exhibit risk aversion when making high-stake decisions (Koenig & Liu 1999). Planning with more realistic planning objectives is an important topic because the recommendations of planners should reflect the opinions of their users correctly. After all, the planners make suggestions for how to act and should make the same suggestions that their users would have made themselves. Utility theory (von Neumann & Morgenstern 1947; Keeney & Raiffa 1976) has studied empirical and normative planning objectives of human decision makers. However, it specifies only what optimal plans are but not how they can be obtained efficiently. Since utility theory is considered to be part of decision theory, decision-theoretic planning, a term often used synonymously with probabilistic planning, is not complete without investigating how to plan efficiently with planning objectives from utility theory. One topic of research in artificial intelligence planning in the past couple of years has been how to plan efficiently with more realistic world models than was possible before, utilizing the structure of the planning problems. The question is whether this structure also allows planners to plan efficiently with more realistic planning objectives. The planning competition last year has made a first small step in this direction and there is some on-going research but it seems that there is now less effort on the topic than there used to be, and a lot of this work is in the UAI (Uncertainty in Artificial Intelligence) community rather than the ICAPS community.

Re-Planning

Planners from artificial intelligence are often one-shot planners. However, many artificial intelligence systems have to adapt their plans continuously to changes in the world or changes of their models of the world. In these cases, the original plans might no longer apply or might no longer be good. Thus, they need to re-plan for the new situations (desJardins *et al.* 1999). Examples of practical significance include the aeromedical evacuation of injured people in crisis situations (Kott, Saks, & Mercer 1999) and air campaign planning (Myers 1999). Similarly, they need to solve series of similar planning problems if one wants to perform series of what-if analyses or if the cost of planning operators, their preconditions, or their effects change over time because they are learned or refined. Consequently, planning is often a repetitive process. In these situations, it can be inefficient to re-plan from scratch, that is, solve the planning problems independently. Fortunately, the changes to the planning problems are usually small. For example, planes might no longer be able to land on a particular airfield for the aeromedical evacuation example. This suggests that some of the previous planning results can be reused to speed up re-planning. There used to be lots of research in this direction in artificial intelligence, where one even distinguished between re-planning and plan re-use. Re-planning attempts to retain as many plan steps of the previous plan as possible whereas plan re-use does not have this requirement. Re-planning and plan re-use methods, for example, include case-based planning, planning by analogy, plan adaptation, transformational planning, planning by solution replay, repair-based planning, and learning search-control knowledge. These re-planning methods have been used as part of systems such as CHEF (Hammond 1990), GORDIUS (Simmons 1988), LS-ADJUST-PLAN (Gerevini & Serina 2000), MRL (Koehler 1994), NoLimit (Veloso 1994), PLEXUS (Alterman 1988), PRIAR (Kambhampati & Hendler 1992), SPA (Hanks & Weld 1995), and SHERPA (Koenig, Furcy, & Bauer 2002). NoLimit, for example, accelerates a backward-chaining nonlinear planner that uses means-ends analysis, SPA accelerates a causal-link partial-order planner, PRIAR accelerates a hierarchical nonlinear planner, and LS-ADJUST-PLAN accelerates a planner that uses planning graphs. There continues to be some research in this direction but it seems that there is now much less effort on the topic in the ICAPS community than there used to be.

References

Alterman, R. 1988. Adaptive planning. *Cognitive Science* 12(3):393–421.

Dean, T.; Firby, J.; and Miller, D. 1988. Hierarchical planning involving deadlines, travel times, and resources. *Computational Intelligence* 4(4):381–398.

desJardins, M.; Durfee, E.; Ortiz, C.; and Wolverton, M. 1999. A survey of research in distributed, continual planning. *Artificial Intelligence Magazine* 20(4):13–22.

Drabble, B.; Koehler, J.; and Refanidis, I., eds. 2002. *Proceedings of the AIPS-02 Workshop on Planning and Scheduling with Multiple Criteria*.

Gerevini, A., and Serina, I. 2000. Fast plan adaptation through planning graphs: Local and systematic search techniques. In *Proceedings of the International Conference on Artificial Intelligence Planning and Scheduling*, 112–121.

Haddawy, P., and Hanks, S. 1992. Representation for decision-theoretic planning: Utility functions for deadline goals. In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning*.

Hammond, K. 1990. Explaining and repairing plans that fail. *Artificial Intelligence* 45:173–228.

Hanks, S., and Weld, D. 1995. A domain-independent algorithm for plan adaptation. *Journal of Artificial Intelligence Research* 2:319–360.

Kambhampati, S., and Hendler, J. 1992. A validation-structure-based theory of plan modification and reuse. *Artificial Intelligence* 55:193–258.

Keeney, R., and Raiffa, H. 1976. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. Wiley.

Koehler, J. 1994. Flexible plan reuse in a formal framework. In Bäckström, C., and Sandewall, E., eds., *Current Trends in AI Planning*. IOS Press. 171–184.

Koenig, S., and Liu, Y. 1999. Sensor planning with nonlinear utility functions. In *Proceedings of the European Conference on Planning*.

Koenig, S.; Furcy, D.; and Bauer, C. 2002. Heuristic search-based replanning. In *Proceedings of the International Conference on Planning and Scheduling*.

Kott, A.; Saks, V.; and Mercer, A. 1999. A new technique enables dynamic replanning and rescheduling of aeromedical evacuation. *Artificial Intelligence Magazine* 20(1):43–53.

Myers, K. 1999. CPEF: A continuous planning and execution framework. *Artificial Intelligence Magazine* 20(4):63–69.

Simmons, R. 1988. A theory of debugging plans and interpretations. In *Proceedings of the National Conference on Artificial Intelligence*, 94–99.

Veloso, M. 1994. *Planning and Learning by Analogical Reasoning*. Springer.

von Neumann, J., and Morgenstern, O. 1947. *Theory of games and economic behavior*. Princeton University Press, second edition.