

OVERCOMING LIMITATIONS OF ARTIFICIAL INTELLIGENCE PLANNING TECHNIQUES

Curtis P. Langlotz, Lawrence M. Fagan, Edward H. Shortliffe

Medical Computer Science Group
Knowledge Systems Laboratory
Stanford University, Stanford, California

Many attempts have been made to provide automated assistance for planning problems. Such research in artificial intelligence (AI) has largely concentrated on solving two kinds of planning problems: *State space search* is used to solve problems for which there is certainty about the consequences of action and for which the planning goals can be met completely. *Skeletal plan refinement* is used to solve problems for which there exist explicit guidelines for the construction of plans. However, many planning problems are characterized by a lack of explicit plan construction guidelines, goals that are difficult to satisfy completely, and actions whose consequences cannot be predicted with certainty. A cancer chemotherapy planning problem is used to illustrate the limitations of these two artificial intelligence techniques, and a new approach is proposed which explicitly represents the uncertainty and tradeoffs inherent in many planning problems.

Introduction

The generic planning process involves determining an appropriate course of action that conforms as well as possible to a set of explicit goals. Fig. 1 illustrates the typical structure of artificial intelligence planning programs. The planner typically has some representation of the current state of the world, and the desired state or goal state of the world. The planner is provided with a set of possible actions or *plan operators* that can be carried out, together with the domain knowledge needed to select operators and predict their effects. Planning strategies may be used to aid the search for an appropriate action. Ultimately, the planner provides a set of operators which, when carried out, will transform the initial state into the goal state.

In the sections that follow, we have identified two AI approaches to planning: state space search and skeletal plan refinement. The analysis highlights the assumptions made by these techniques for planning under uncertainty, and indicates the need for a new architecture. The approach that has emerged from our current work to date is then outlined.

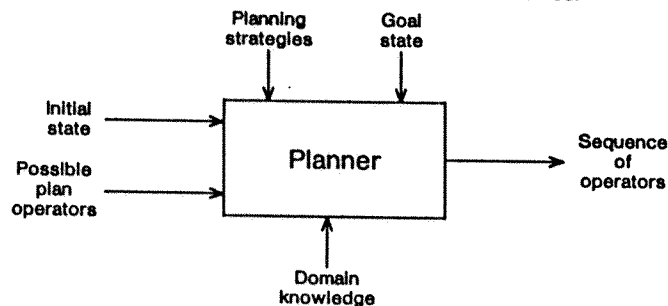


Figure 1: The generic architecture of a typical artificial intelligence planning system.

AI Planning Using State Space Search

State space search systems typically represent both the initial state and the goal state as predicate calculus assertions (1). Plan operators are represented in terms of: (a) their preconditions (indicating which assertions must be true of any state to which an operator can apply) and (b) their effects (indicating which assertions must be added to or deleted from a state each time an operator is applied).

In this model, the planner searches through the space of all possible states using plan

Overcoming limitations of AI planning techniques

operators to move between states. A variety of planning strategies are used to aid the search for a path leading from the initial state to the goal state. For example, means-ends analysis (2) recommends those operators which will most reduce the difference between the state under consideration and a state which satisfies the goal. More advanced planners have used strategies that deal with the problem of interacting goals and plan operators. NOAH (3) uses a critic to repair incorrect or inconsistent plans for block stacking problems. Stefik's MOLGEN (4) uses constraints and a least commitment approach to help avoid overcommitting to a particular plan in the domain of molecular genetics experiment design. All these planners provide a sequence of operators designed to satisfy the goal.

State space search planners are effective in finding innovative orders of operators to satisfy planning goals. However, systems of this kind make a number of assumptions:

- the effects of plan operators must be stated explicitly and with certainty;
- the goals of the planning problem must be stated explicitly and with certainty;
- the goals must be stated in such a way that they can be completely satisfied (i.e., unresolvable tradeoffs must not occur between the parts of the goal).

These assumptions limit the applicability of such planning techniques in many areas. Consequently, most planning systems of this type have solved simple block stacking problems or robot movement problems.

AI Planning by Refining Skeletal Plans

For many planning problems, the plan operators can be grouped into classes which can apply at specific points in the problem solving process. This sequence of classes is called a *skeletal plan*. Domain knowledge and strategic knowledge are used to *refine* the operators in the skeletal plan to ones which are appropriate for the problem at hand. Some planning programs which rely on this skeletal plan refinement approach are: a parallel effort to Stefik's molecular genetics experiment planner (also called MOLGEN) developed by Friedland (5), VM, which assists with the management of mechanical ventilation of post-surgical patients (6), ONCOCIN, which assists physicians with the management of chemotherapy of cancer patients (7), and ATTENDING which critiques physicians' anesthetic management plans (8). Because a skeletal plan is provided to these systems in advance, they construct plans which have a consistent overall structure. For example, ONCOCIN is not concerned with *when* in the plan sequence to give drug therapy, but rather which of a number of possible variations of drug therapy should be given. Similarly, ATTENDING is not concerned with *when* in the plan sequence to induce anesthesia in a patient, but rather with choosing the most appropriate induction method from the ones which are available.

Both ONCOCIN and ATTENDING represent their skeletal plans as a constant hierarchical structure. In ONCOCIN, this has been called the *recommendation hierarchy* (9); in ATTENDING, it has been called the *physician approach tree* (10). Both these systems use additional domain knowledge to refine the plan steps at each point in the recommendation hierarchy. For example, the rules shown in Fig. 2 are used in ONCOCIN to refine generic planning steps such as "order tests" and "give drugs" to more specific actions.

```

IF: The chemotherapy that the patient is expected to receive on the
next visit is HD-MTX
THEN: Recommend that the Serum Creatinine test should be done On Next
Visit STAT
  
```

```

IF: 1) The patient has received chemotherapy, and
2) The measure of neurologic toxicity is 1
THEN: The dose of Vincristine to be given is 50 percent of the
100% dose.
  
```

Figure 2: Paraphrased versions of two rules used by ONCOCIN in deciding which actions to perform when planning cancer therapy. The first rule is used to determine which test to order. The second is used to decide on precise dose attenuations.

Skeletal planning is useful when expert guidelines for selecting plan operators can be readily expressed and innovative plan orderings need not be generated (5). Planning systems that use this methodology assume that an explicit plan order can be specified in advance and that, accordingly, planning goals remain relatively constant between problems.

Because a skeletal plan is refined according to the characteristics of specific planning situations, this technique only yields a solution when the system builder can anticipate the characteristics of most planning problems which are likely to occur. Therefore, it is not feasible for this kind of planning system to provide planning solutions for every possible situation in a complex domain; this would involve specifying knowledge about how to respond to an extraordinary number of specific planning situations.

Overcoming limitations of AI planning techniques

The Difficulties in Planning with Uncertainty and Tradeoffs

ONCOCIN (7) is a skeletal planning system that encodes the knowledge contained in cancer treatment protocols together with experts' interpretations of those protocols. This knowledge is encoded in rules which link specific treatment situations to treatment options. Published studies have shown that ONCOCIN offers advice which is comparable to the treatment offered by that at a university oncology clinic (11). The data suggest that the program already has the knowledge to provide useful assistance for many patients with complex oncologic problems. However, ONCOCIN (and the protocol that describes a specific experimental treatment regimen) cannot possibly anticipate all potential problems with therapy, nor can it list guidelines for all appropriate responses. ONCOCIN accordingly refers patients who are not covered by the protocol guidelines to the study chairman, who then recommends therapy using his or her knowledge about the physiology of the human body and the disease process, together with strategies for oncology chemotherapy and the goals for the patient and the clinical study.

We wanted ONCOCIN to be able to provide planning assistance even for those cases which might otherwise be referred to the study chairman. This led us to consider either enhancing our existing skeletal planning methodology or using state space search to augment ONCOCIN's advisory capabilities. However, we discovered that the assumptions made by state space search cannot be met for our problem and that explicit skeletal planning guidelines were impractical to devise for such a wide range of possible planning situations.

We set out to design a computer program, called ONYX, that could handle these more difficult planning problems. Fig. 3 shows the design of ONYX, viewed as an analogy to the general planning process outlined in Fig. 1.

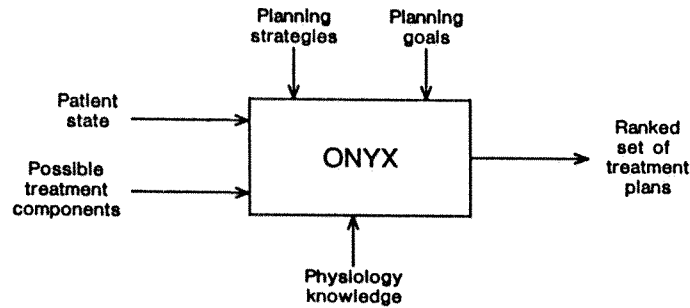


Figure 3: The structure of ONYX when considered within the framework of previous AI planning systems.

The initial state in this new planning problem is the patient's state at the time of therapy. Unfortunately, the description of that state is partially dependent on (a) the temporal trends of the patient's response to previous therapy, and (b) variables which are not directly measurable. For example, relative numbers of various precursor cells in the patient's bone marrow may be crucial to an assessment of the patient state, but it is not practical to examine the marrow each time therapy is given. These considerations imply that substantial uncertainty about the current state of the patient must remain at the time therapy is planned.

Oncology experts use knowledge about the behavior and structure of the human body in the planning process to help predict the possible consequences of plan steps. Although the structure of the body is relatively invariant between patients, there is substantial variability in how patients will behave in response to the same therapeutic intervention. The uncertainty about the behavior of the body implies that there must also be some uncertainty about the consequences of carrying out certain plan operators. A dose of a given drug, for example, may have widely varying effects in patients. Furthermore the effects of the drug may occur with some delay, may be transient, or may not be detectable with certainty.

There are clear goals to be met in oncology therapy planning: "Kill the tumor cells", "Avoid toxicity". But these goals, unlike the goals in traditional planners, are not completely satisfiable. The planner must necessarily make tradeoffs between them. For example, the larger the drug dose given, the less that the tumor will grow (a good thing), but the more that the patient will experience toxicity (a bad thing).

Because of the uncertainty involved, any proposed set of operators cannot be guaranteed to transform the initial state into the goal state; instead the set of operators should yield the maximum expected attainment of goals.

It is evident that planning problems of this kind do not meet the assumptions required by the state space search and skeletal planning methodologies. The following statements summarize the salient features of this problem that distinguish it from other planning problems:

- explicit guidelines for plan selection are not available;
- there is uncertainty about the consequences of action;
- planning goals cannot be satisfied completely;
- the initial state is not known with certainty.

An Architecture for Planning with Uncertainty and Tradeoffs

It has become clear to us with experimentation that in the ONCOCIN domain it is important to represent explicitly the uncertainties about patients' responses to interventions and the tradeoffs between planning goals. We also have found ourselves unexpectedly but (in retrospect) inevitably drawn to the possible use of established normative techniques, namely the decision theoretic approach for deciding among a group of alternative plans (12). This theory reduces the AI planning problem to three subproblems: (i) to generate the group of alternate plans which should be considered in the decision analysis, (ii) to determine the probabilities of possible outcomes if the competing potential plans were actually carried out, and (iii) to represent the utility of possible outcomes and adjust them for individual patients.

The desire to exploit this new problem structure and its concomitant advantages for decision-making and explanation were strong motivations in the design of a new planning approach. This new planning architecture is shown in Fig. 4.

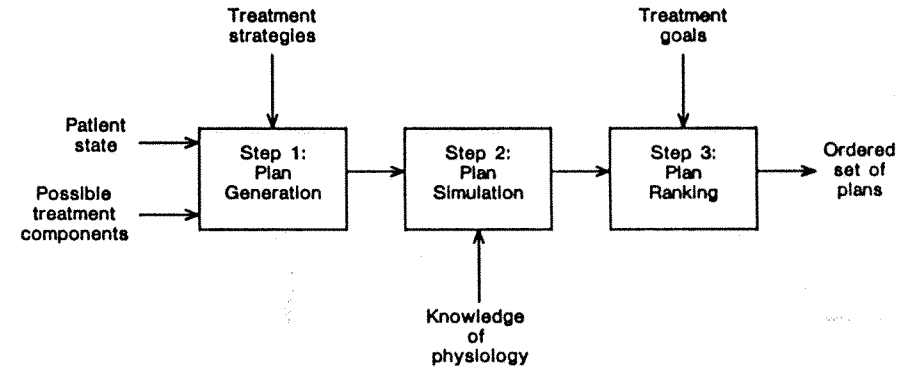


Figure 4: The three step process for planning under uncertainty in ONYX.

As shown in the figure, the planning process consists of three steps:

- Plan generation.** Using oncology treatment strategies, generate a small set of reasonable plans by selecting combinations of treatment components appropriate for the current patient state.
- Plan prediction.** Using knowledge about the structure and behavior of the human body, together with probabilistic models of human physiology, predict (through simulation) the future states of the patient after the execution of each proposed plan.
- Plan ranking.** Using decision analysis, rank the plans according to how well the predictions for each plan meet the therapeutic goals for the patient. The highest ranked plan can be selected for action.

Our experience with each of these components has led to an understanding of this model for decision-making; we believe it will lead to a computer program which can provide intelligent assistance in planning with uncertainty. Our current work on this problem is embodied in ONYX, the planning system designed to be called by ONCOCIN when a patient fails to meet the evaluation criteria delineated in a formal protocol. ONYX has been implemented in prototype form on a Xerox 1108 Lisp machine, and it explicitly adheres to the 3-step process outlined in Fig. 4.

Conclusion

Uncertainties and tradeoffs are pertinent in a wide variety of planning problems analogous to those we are encountering in augmenting ONCOCIN's advisory capabilities. We were surprised to find that AI techniques had substantial limitations when solving planning problems of this kind. We believe that the essential features of the new architecture we have devised, in its attempt to explicitly capture uncertainties and tradeoffs, will be a useful in a wide variety of planning

Overcoming limitations of AI planning techniques

situations where automated assistance might not otherwise be possible.

Computer programs which provide assistance with medical planning problems almost inevitably require explicit representation of uncertainty. Probability theory, decision analysis, and simulation are established techniques which are effective in solving problems under uncertainty. However, these techniques only apply after the problem has been properly structured. We believe that AI techniques can be used to help structure problems for these normative techniques: for example, by generating a set of alternative plans for decision analysis, providing the probability of outcomes for decision analysis, tailoring a mathematical simulation to correspond to the known past characteristics of an individual patient, or modifying a utility model based on patient characteristics. The planning problems we are addressing indicate the continuing need to combine established normative techniques with the symbolic reasoning and representation techniques developed in artificial intelligence research. We believe that these combinations are fruitful ones and will likely lead to enhanced decision support for problems where uncertainty and tradeoffs are dominant problem features.

Acknowledgements

Support for this work was provided by the National Library of Medicine under Grants LM/CA-04136 and LM-04318, the National Science Foundation under Grant IST83-12148, the Division of Research Resources under Grant RR-01636, and Ford Aerospace under Grant SP882573CJ. Computing facilities were provided by the SUMEX-AIM resource under NIH Grant RR-00785 and by an equipment grant from the Xerox Corporation. We appreciate the continued collaboration of the members of the ONYX project: Mark Frisøe, Michael Kahn, Hank Rappaport, Branimir Sikic, Samson Tu, and Cliff Wulfman.

References

1. Fikes, R. E., and Nilsson, N. J. "STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving." *Artificial Intelligence* 2 (1971), 189-208.
2. Newell, A., and Simon, H. GPS, A Program That Simulates Human Thought. In *Computers and Thought*, McGraw-Hill, New York, 1963.
3. Sacerdoti, E. D. *A Structure for Plans and Behavior*. Elsevier, New York, 1977.
4. Steilk, M. J. "Planning With Constraints (MOLGEN: Part 1)." *Artificial Intelligence* 16 (1981), 111-139.
5. Friedland, P. E. *Knowledge-based experiment design in molecular genetics*. Ph.D. Th., Computer Science Department, Stanford University, 1979. Rep. No. STAN-CS-79-771
6. Fagan, L. M., Kunz, J. C., Feigenbaum, E. A., and Osborn, J. J. Representation of dynamic clinical knowledge: Measurement interpretation in the intensive care unit. Proceedings of the 6th International Joint Conferences on Artificial Intelligence, Tokyo, 1979, pp. 260-262.
7. Shortliffe, E. H., Scott, A. C., Bischoff, M., Campbell A. B., van Melle, W., Jacobs, C. ONCOGIN: An Expert System for Oncology Protocol Management. Proceedings of the Seventh International Joint Conference on Artificial Intelligence, Vancouver, British Columbia, 1981, pp. 876-881.
8. Miller, P. L. "Critiquing Anesthetic Management: the ATTENDING Computer System." *Anesthesiology* 53 (1983), 362-369.
9. Langlotz C. P. and Shortliffe, E. H. "Adapting a Consultation System to Critique User Plans." *International Journal of Man-Machine Studies* 19, 5 (1983), 479-496.
10. Miller, P. L. "ATTENDING: Critiquing a Physician's Management Plan." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 5 (1983), 449-461.
11. Hickam, D. H., Shortliffe, E. H., Bischoff, M. B., Scott, A. C., Jacobs, C. D. "A study of the treatment advice of a computer-based cancer chemotherapy protocol advisor." *Annals of Internal Medicine* (1985). (In press.)
12. Raiffa, H. *Decision Analysis: Introductory Lectures on Choice Under Uncertainty*. Addison-Wesley, Reading, Mass., 1968.