Challenges of Decision-Theoretic Planning

Michael P. Wellman
University of Michigan AI Laboratory
1101 Beal Avenue, Ann Arbor MI 48109
wellman@engin.umich.edu

1.0 Introduction

Several scientific and engineering disciplines study techniques for automated decision making. For example, work in control theory, operations research, and management science has produced numerous automated or semi-automated methods for deriving prescribed actions from descriptions of situations and models of the world. The type of techniques studied in these fields is shaped by the primary problems of interest, predominant technical methodologies, and available modeling technologies. Similarly, the shape of research in *planning*–the decision-oriented discipline of artificial intelligence–has been molded by concerns specific to AI research, and strongly influenced by its technical culture and developmental history.

Perceived weaknesses in the capabilities of AI planning systems and fundamental inadequacies in the basic planning framework are sometimes contrasted with perceived strengths and adequacies in other decision disciplines. Invariably, the advantages accrued elsewhere come at the cost of some other restrictions or limitations AI researchers would find otiose. In attempting to unify the best aspects of several approaches, it is necessary to understand the underlying sources of their respective strengths and weaknesses. Interactions among these sources often pose obstacles to combining concepts and techniques.

Recent years have seen a substantial increase (from nil to non-negligible) in interest among the AI planning community in concepts from the economic theory of rational decision making, or *decision theory*. Decision theory underlies several approaches to automated decision making and decision support, notably the applied methodology of *decision analysis* studied in operations research and management science. Many are persuaded by the theory's claim to characterize an idealized form of rationality, finding it attractive for this reason as well as its generality. Nevertheless, neither decision theory nor decision analysis *per se* address all the concerns of AI planning, and hence those who find decision-theoretic principles appealing must find a way to realize this notion of rationality within a plan-synthesizing computational agent. In pursuing this goal, one might expect (or at least hope) that existing ideas from AI planning will play a useful role, leading to a computational decision theory unifying the two disciplines.

Let us apply the term *decision-theoretic planning* to the enterprise of incorporating decision-theoretic principles into AI planning, both theory and practice. In this paper, I define this enterprise more precisely and outline the major technical challenges. This short exposition is not meant to survey current research relevant to the enterprise (or even cite any of it), but rather to frame the problem and explore some of the central issues facing researchers in this emerging area of investigation.

2.0 Definition

In brief, decision theory takes as primitive a set of plans, for each plan a probability distribution over its possible outcomes, and a utility function over outcomes. The recommended plan is that maximizing expected utility. To define what it means to apply these concepts to AI planning, we need first to identify the

essential aspects of planning. Abstracting from some of the details of common AI planning languages, we can derive a general definition of the task:

> Given (1) a schematic description of primitive actions and their effects, and (2) a description of objectives, the *AI planning task* is to (3) synthesize a plan serving those objectives, where a *plan* is a compound object consisting of action instances connected by basic relational primitives.

The basic relational primitives are composition operators, such as sequencing or conditioning, that specify how the individual actions are related in the composite plan object. Note that this definition is far more general than typical characterizations of AI planning, which tend to incorporate features of particular planning approaches. For example, we might define *classical planning* as a specialization of the above where objectives are described by goal propositions and sequencing is the basic composition operator. Other specializations might constrain the representation of effects (e.g., by preconditions, add-, and delete-lists) or the intermediate structures explored during plan synthesis (e.g., linear or nonlinear plan fragments). Whereas distinctions such as these are technically significant and useful, they are not definitive for AI planning. Rather, the essential characteristic of AI planning that distinguishes it from other decision disciplines is its focus on structured plan objects composed from schematically described action primitives. This leads to a compact but expressive description of a combinatorial but highly structured plan space. As long as the effects of actions compose in a modular manner, this regularity can be exploited in plan synthesis.

As for classical planning, we can define decision-theoretic planning via specialization:

> *Decision-theoretic planning* is a special case of the general AI planning task where (1) effects of actions are described by probability distributions over outcomes, (2) objectives are described by utility functions, and (3) the criterion for effective plan synthesis is expected-utility maximization.

By "described by" in the above, we do not mean to insist that the languages for effects and objectives literally and directly represent complete probability distributions and utility functions. Rather, we require only that the representation language have some interpretation in terms of those concepts, such as constraining them in some way.

Seen as sibling specializations of a general planning task, the most critical differences between classical and decision-theoretic planning are that the latter admits explicit representation of uncertainty in the effects of actions and graded preferences for alternate outcomes. These extended specifications provide a sharper rationality criterion than satisficing, imposing a more refined preference among plans than possible when the objective is merely a goal proposition.[1]

3.0 Challenges

Classical planning and its variants are difficult for two primary reasons: (1) the need to search a combinatorial or infinite plan space (i.e., even restricted planning problems tend to be NP-hard), and (2) a variety of subtle technical problems arise in reasoning about the effects

1. Newell elucidates the rationality theory corresponding to satisficing in his original exposition of knowledge-level analysis. A rationality theory based on goals alone is fundamentally limited in that a goal cannot distinguish among plans that achieve it, nor prescribe behavior when a satisficing plan cannot be found.

of actions (e.g., ramifications, persistence, qualifications, combined or context-dependent effects). These problems are no less serious in decision-theoretic planning, and in some cases they are more difficult. In this section I outline some of the technical challenges specific to the decision-theoretic approach. The magnitude of these difficulties suggests that we should not expect to construct a decision-theoretic planner by some incremental modification to a deterministic-action goal-satisficing planner (e.g., by simply inserting probabilities and utilities somewhere). However, they should not be taken as an argument against the enterprise, as even partial success in accounting for likelihood and preference can significantly extend the applicability of AI planning technology.

It should be noted that some of the difficulties enumerated below can be alleviated or eliminated given sufficient restrictions on the problem domain. Indeed, the literatures of the decision disciplines mentioned above catalog a myriad of special cases. However, AI research has traditionally been concerned with general-purpose mechanisms, and hence it would be inadmissible to hardwire the necessary assumptions into a planning architecture. Instead, we should seek to develop general mechanisms that can exploit regular structures and restricted forms when they exist. It is in this regard that the knowledge and techniques developed by the other decision disciplines will be of greatest use to artificial intelligence.

3.1 Schematic action representations

Perhaps the main feature of AI planning that we must accommodate in adopting the decision-theoretic concepts of probability and utility is that of schematic action representations. Standard representations of probability distributions and utility functions apply to event and outcome *instances*, as opposed to *types*. In other words, the well-developed languages for probabilistic and utility-theoretic reasoning possess propositional expressive power. Although it is relatively easy to obtain a first-order language by writing down the axioms for probability and utility, the challenge is to design a language that supports general expressions about patterns of likelihood and preference with respect to various actions and conditions.

3.2 Uncertain effects of actions

Once we admit that we cannot flawlessly predict the effects of actions, several problems emerge that did not arise in the deterministic case.

First, uncertain relationships are inherently less modular than their deterministic counterparts. Whereas logical connectives are truth functional in the deterministic case–e.g., the truth value of a conjunction is a function of the truth value of the conjuncts–the same does not hold of probabilities. The probability of a conjunction is constrained by the probabilities of its conjuncts (it cannot exceed them), but is not uniquely determined in general. This means that we must consider the interactions among the parts in assessing our belief in the whole. Similarly, whereas in the deterministic case we need only find one successful argument (proof) for a proposition, under uncertainty we must in general find all relevant evidence pro and con and weigh it together in order to establish the proper degree of belief for the proposition. In general, it is more difficult in the uncertain case to achieve separability in representation and reasoning. For this reason, we should expect increased complexity in specifying and reasoning about effects of actions in planning under uncertainty.

Second, maintaining graded degrees of belief rather than binary truth values imposes an added burden of precision. Precision generally increases the complexity of knowledge specification, if only because more bits are required.

Worse, precise relationships tend to be more context sensitive than coarse ones, since features that might not have affected a categorical belief value may nevertheless influence the exact degree of belief. This is especially problematic when we consider the temporal dimension usually present in planning problems. Whereas it is reasonable to consider changes in binary truth values at only discrete time points, degrees of belief typically change more frequently, even continuously. Like non-modularity, the precision factor requires us to consider more and higher-order relationships, thus inhibiting scalability of our planning knowledge.

Third, the inability to perfectly predict effects means that some important information about the world may become available only at execution time. In order to make use of such information, our plans must include conditional constructs enabling action to depend on downstream observations. Execution-time replanning is not sufficient, as (1) it takes time, which may not be available, and (2) the ability to react to observations downstream may impact the preferred actions taken at the beginning of the plan. Accounting for contingencies dramatically blows up the plan space, rendering an intractable problem that much more so.

Finally, the opportunity to condition action on dynamically obtained information suggests that it may be advantageous to expressly plan actions that resolve salient uncertainty. As above, considering information-gathering actions adds another level of combinatorial explosion to an already enormous search space.

Note that these problems are attributable to uncertainty and degrees of belief generally, and are not specific to the probabilistic approach. Some formalisms may appear immune to some of these problems (e.g., logical connectives in the fuzzy calculus *are* truth functional), but only by defining away the potential interactions or opportunities for contingency planning that make the general problem difficult.

3.3 Preferences over outcomes

Accounting for relative preference is perhaps the most fundamental feature of decision-theoretic planning distinct from traditional AI planning approaches. Most AI planning theory and technology presumes that objectives are represented by goal propositions; indeed, this representation is often taken as part of the definition of the planning task. That goal satisfaction is so firmly rooted in AI is probably attributable to two main factors:[2] (1) the prevalence of predicate-logic representations, which naturally capture binary distinctions among states, and (2) the seminal influence of Herbert Simon, whose theory of satisficing behavior challenged the optimizing framework standard in economics.

Generalizing objectives to permit relative preference among states presents two significant challenges. First, we need a language to specify these preferences. And second, given specified preferences, we require an effective procedure for finding plans with preferred outcomes.

One of the central results of decision theory is that, given some conditions on rational choice, the preference order is representable by a real-valued utility function. But this result does not solve our planning representation problem, as numeric utility functions suffer from the same plagues of precision and modularity enumerated for probability above. Plausible representation for relative values must capture the robust patterns of preference among outcomes, and provide for modular combination of mea-

2. warning: reckless historical speculation ahead

sures of distinct attributes of value. Goal propositions constitute one important class of preference pattern for which decision-theoretic planning must account; research to discover the other important classes has barely scratched the surface.

Simon's argument for satisficing as a form of "bounded rationality" was based on empirical investigation of human decision makers as well as the insightful observation that satisfying a preset aspiration level is an easier computational problem than global optimization. Measuring the difference in complexity is difficult outside some particular context, but this assertion can be made precise in specific formulations. One advantage of satisficing is that it is locally testable; in contrast, one cannot generally tell by inspection of an isolated solution that it is optimal on some graded measure. To determine optimality without exhaustive search requires that we exploit regularity in the search space and some knowledge about the regular structure.

3.4 Analyzing partial plans

In some attempts to integrate planning and decision theory, the two ideas are applied serially–plans are generated and then evaluated. Given the size of typical plan spaces, this approach can ultimately evaluate only a minuscule fraction of the possible plans. What basis have we to expect these to include the good ones? The criterion for plan quality, after all, is specified in the decision-theoretic part. In order to find good plans, we must bring in some decision-theoretic evaluation before plans are fully assembled. This is not completely straightforward, as the standard concepts are not defined for partial plans.

The difficulty in assessing the value of a partial plan is that there are many ways in which that plan may be completed. In the usual case, the set of completions include some that are arbitrarily bad, and some that are quite good. Comparison of partial plans must be based on some indications of which partial plans have the best completions, or perhaps those best and easiest to find.

It is unclear whether in decision-theoretic planning we can define useful analogs to the sort of modal truth criteria applied in classical planning. That is, whether we can effectively assess properties (even uncertain ones) holding at a point of a partially specified plan. The inherent non-modularity of uncertain relationships makes this more difficult than the classical case, but it is still too early to tell whether it is a reasonable approach.

4.0 Conclusion

The technical challenges facing the enterprise of decision-theoretic planning are deep and numerous. In this paper, I have not attempted to outline the possible strategies for meeting these challenges, nor to describe the ideas that have already been advanced for overcoming or alleviating some of the specific difficulties addressed. Such a survey will find another occasion. For the nonce, suffice it to say that several AI researchers have undertaken some of these challenges in recent years, and partial solutions are emerging. At the moment it is too early for me to assess with confidence the prospects for a general-purpose decision-theoretic planner truly worthy of the name. Nevertheless, I maintain a high degree of confidence that the enterprise is worthy of effort, as (1) the ultimate goal is perhaps the closest characterization of "an ideally rational agent", which should be a holy grail for AI, (2) even partial success could dramatically extend the scope of AI planning and increase the competence of planning systems, and (3) we haven't yet tried hard enough to give up.