Editorial

Planning in artificial intelligence is concerned with developing methods that reason about the consequences of actions to determine how best to achieve given objectives. Planning is a key technology for building intelligent systems in areas such as manufacturing, space systems, software engineering, robotics, education, and entertainment.

Planning, following in the footsteps of automatic control, assumes a set of possible states and a set of possible actions, though typically these sets are discrete and finite. Classical planning methods usually assume that the start state is completely known and that each action corresponds to a function that deterministically transforms any state to a successor state; the successor state depends only on the action and previous state. The planning objective is to find a short or minimal-cost action sequence that transforms the start state to a state that satisfies a goal criterion. Thus, planning problems can be viewed as graph search problems whose vertices are states and whose edges are actions. However, since states are typically described in terms of the presence or absence of features and their number is therefore exponential in the number of features, it is infeasible to use conventional graph search algorithms for planning.

Researchers have therefore studied how to exploit structure to represent planning problems concisely (often using a particular problem representation called STRIPS) and solve them efficiently. For example, it is possible in some cases to take advantage of regularities in how actions affect particular features (only a small number of features may be affected) to produce compact representations of problems with very large state spaces. Planning with STRIPS representations is PSPACE complete. Despite this limitation, a lot of progress has been made towards solving larger and larger planning problems. Recently, novel planning techniques have been developed by integrating technology from other areas into planners. One example involves using binary decision diagrams from symbolic model checking to represent large sets of states compactly and manipulate them efficiently. It is also possible to interpret planning problems as boolean satisfiability (SAT) problems and use fast SAT-solvers to tackle them.

Classical planning does not account for uncertainty. However, planning in the real world has to deal with uncertainty and, furthermore, must find solutions in a timely manner despite the large number of resulting contingencies. In crisis situations such as marine oil spills, for example, planners face uncertainty about how the oil drifts, how the weather conditions change, how long it takes vessels to reach the spill, and so on, yet they have to determine how to contain the oil before it reaches the shores. Similarly, mobile robots have to deal with uncertain actions (actuator uncertainty) and uncertain observations of their state (sensor uncertainty) while coordinating their behavior with a dynamic environment and carrying out tasks in a timely manner. Uncertainty typically explodes the number of contingencies a planner has to contend with.

There are different ways of relaxing the assumptions of classical planning to model uncertainty:

- First, one has to decide which kinds of uncertainty to model, for example, uncertainty about the start state or about the effects of actions. In the presence of such uncertainty, it is no longer possible to predict with certainty which states will result from action executions. Thus, it becomes important to determine the current state by sensing when possible. Sensors are often noisy, however, and so in many cases the planner will remain uncertain about some aspects of its current state.
- Second, one has to decide on how to model the uncertainty, for example, symbolically using sets (say, by enumerating all states that can result from an action execution) or numerically using probability distributions over these sets.
- Third, one has to decide which plans to consider, for example, closed-loop ("contingent") plans or open-loop ("conformant") plans, that is, action sequences.
- Fourth, one has to decide on the planning objective. If uncertainty is represented symbolically, the planning objective could be to find a plan that has some possibility of achieving the goal (weak planning) or one that is guaranteed to achieve the goal (strong planning). If uncertainty is represented with probabilities, the planning objective could be to find a plan that maximizes the probability of achieving the goal or one that minimizes the expected cost until the goal is achieved.

This listing shows that there are a variety of different assumptions one can make. These assumptions differ in which planning problems they apply to, how quickly the planning problems can be solved, and which planning methods are appropriate for solving them. This variability in basic assumptions explains the diversity of approaches for planning under uncertainty.

If uncertainty is represented probabilistically, then goal-directed versions of Markov decision processes from operations research provide a foundation for planning under uncertainty since many such planning problems can be modeled using this framework (just like many classical planning problems can be modeled as graph search problems). Fully observable Markov decision processes (MDPs) generalize graphs to model actuator uncertainty, whereas partially observable Markov decision processes (POMDPs) generalize graphs to model both actuator and sensor uncertainty. The probability of a successor state for Markov decision processes depends only on the current state and the executed action (Markov assumption), similar to classical planning. Operations research has developed dynamic programming methods (such as value iteration and policy iteration) that determine closed-loop plans that typically minimize the expected execution cost over a finite or infinite planning horizon. Thus, if uncertainty is represented probabilistically, one can, in principle, use methods from operations research for planning under uncertainty. Similarly, if uncertainty is represented symbolically, one can apply AND-OR graph search methods to the planning problem. As in the case of classical planning, however, large state spaces make it infeasible to use conventional algorithms for planning. Borrowing insights from classical planning, researchers use regularities of the action effects, now in the form of conditional independence involving features of the states, to achieve compact representation and efficient inference.

This special issue of Artificial Intelligence on "Planning with Uncertainty and Incomplete Information" documents some of the recent progress in this exciting area of artificial intelligence. All seven papers study particular points in the space spanned by two features, namely the expressiveness of the problem representations and the efficiency of solving planning problems. Some authors do this by deriving theoretical results about planning with respect to particular problem representations or planning methods. Most authors, however, do this by deciding on a problem representation, developing planners for that problem representation, and comparing them empirically against the state of the art.

Some authors extend ideas from classical planning to planning under uncertainty, for example, using binary decision diagrams to represent sets of states or interpreting planning problems as stochastic satisfiability problems. Others extend ideas from operations research, for example, to represent and solve MDPs with additional structure, such as factored MDPs. Still others draw on insights from automata theory, concurrent process semantics, and probabilistic finite state automata.

Ultimately, we believe that the key to making progress on planning under uncertainty is to combine ideas from a variety of disciplines that are concerned with how to make good decisions under uncertainty, including artificial intelligence, operations research, decision theory, control theory, and economics. Effective syntheses require serious technical advances to reconcile the different assumptions and approaches. The collection of papers in this special issue shows a representative sample of research in this direction.

The paper "On the Undecidability of Probabilistic Planning and Related Stochastic Optimization Problems" by Madani et al. studies a variety of probabilistic planning problems with sensor uncertainty and infinite planning horizons and shows that they are undecidable, thereby refining our understanding of the hardness of planning under uncertainty.

The paper “Weak, Strong and Strong Cyclic Planning via Symbolic Model Checking” by Cimatti et al. formalizes planning (including the notions of weak and strong planning) where the uncertainty is represented symbolically and develops planning methods that can solve them efficiently with binary decision diagrams.

The paper “SAT-Based Planing in Complex Domains: Concurrency, Constraints and Nondeterminism” by Castellini and Giunchiglia formalizes planning where the uncertainty is represented symbolically and develops planning methods that can solve them efficiently with boolean satisfiability methods. “Contingent Planning under Uncertainty via Stochastic Satisfiability” by Majercik and Littman does the equivalent for probabilistic planning.

The papers “Equivalence Notions and Model Minimization in Markov Decision Processes” by Givan et al. and “Solving Factored MDPs using Non-Homogeneous Partitions” by Kim and Dean formalize planning with factored MDPs and develop efficient planning methods by aggregating “similar” states for different notions of similarity.

Finally, the paper “Performance Bounds for Planning in Unknown Terrain” by Koenig and Tovey formalizes planning for robot navigation in unknown terrain where the uncertainty is represented symbolically and analyzes existing planning methods that have successfully been demonstrated on robots, showing that they are indeed efficient. This paper differs from the other six papers in that it does not solve the planning problem off-line but rather interleaves planning and execution to make planning under uncertainty tractable.

Craig Boutilier (University of Toronto)
Tom Dean (Brown University)
Sven Koenig (Georgia Institute of Technology)