Probabilistic Knowledge of External Events in Planning

Jim Blythe

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213
blythe@cs.cmu.edu

My research tries to improve the robustness of plans by using limited knowledge about external events. These are events that are not directly caused by the planning agent. I use a discrete-time model and assume that the probability of occurrence for a particular type of event in a given situation is known, but the specific occurrence of such an event cannot be predicted with certainty. For example, when a bicycle is left outside a building, there is some probability p that it will be stolen at each time point. The probability that the bicycle is still outside the building after n time units is then $(1 - p)^n$, neglecting the effects of other possible events.

"Reactive" planners, that create plans in response to exceptions during execution, are subject to severe time constraints that limit their performance. My aim is to create contingent plans off-line that have a higher probability of success than could be achieved under these constraints and without considering external events. Stochastic techniques such as policy iteration are not yet tractable for realistic planning problems, so I concentrate on classical planning systems.

Most other planning systems that reason about uncertainty deal with non-deterministic effects and/or incomplete knowledge of the state of the world, e.g. (Kushmerick, Hanks, & Weld 1993; Pryor & Collins 1993). However, these models cannot easily be used to represent uncertainty about external events. There are two main reasons for this, both related to the frame problem and stemming from the fact that uncertainty about events depends on the actions being performed only through the state. Firstly, uncertainty about the occurrence of an event while an action is performed would have to be modelled by probabilistic effects for every action that could be performed at the same time as the event. For instance, every action that can be performed while the bicycle is outside would have to include the bicycle being stolen as a possible effect. Thus, the number of distinct possible outcomes for each action is exponential in the number of event types that could simultaneously occur. This is a misleading representation, as well as very cumbersome. Secondly, consider a planning agent that leaves a bicycle outside a building, performs some actions inside the building and then returns to cycle away. The probability that the bicycle is outside the building, as required, upon the agent's return depends on the length of time spent inside the building, but this fact could not be modelled in the situation calculus using non-deterministic actions alone.

I have designed a planner based on Prodigy 4.0 that considers external events in order to increase the expected utility of a plan, and applied it to a transportation domain. I represent events in a STRIPS-like fashion, with preconditions specifying when events are possible, add and delete lists specifying their effects, and with a probability attached to each event (Blythe 1994). When the preconditions are satisfied, the event may occur with the given probability. The planner first produces a plan without considering external events, and may then use three different routines to find possible sequences of events that would cause the plan to fail: a Monte Carlo simulation, a Markov chain analysis and an exhaustive search for single event instances that can defeat the plan. There is no exhaustive search for sequences of events that would defeat the plan, as this would in general take far longer than the planning phase.

Once sequences of events are found that can cause the plan to fail, the system considers three different ways to repair each one: (1) Conditional steps can be added that address the "bad" situations that arise, for example hailing a cab if the bicycle is stolen. (2) Steps can be added to make bad events inapplicable or less likely, for example by locking the bicycle. (3) The existing steps can be re-ordered or moved along a timeline to reduce the probability of bad events, for example re-ordering the plan to spend as little time as possible inside the building. Since some events have zero duration, this method can sometimes eliminate problems entirely. In experiments, these techniques improve the probability of plan success greatly. I plan to improve plan projection, using techniques such as Bayes nets to relax independence assumptions, and study various domains. I also aim to prove convergence of the algorithm to plans of maximal utility.

References

Blythe, J. 1994. Decision-theoretic subgoaling in goal-directed search. AAAI Spring Symposium on Decision-Theoretic Planning.

Kushmerick, N.; Hanks, S.; and Weld, D. 1993. An algorithm for probabilistic planning. Technical Report 93-06-03, Department of Computer Science and Engineering, University of Washington.

Pryor, L., and Collins, G. 1993. Cassandra: Planning for contingencies. Technical Report 41, The Institute for the Learning Sciences.