

Information Gathering as a Planning Task

Louise Pryor and Gregg Collins*

The Institute for the Learning Sciences
Northwestern University
1890 Maple Avenue, Evanston IL 60201
pryor@ils.nwu.edu, collins@ils.nwu.edu

Abstract

Existing planners fall into two broad categories: reactive planners that can react quickly to changes in the world, but do not project the expected results of a proposed sequence of actions, and classical planners that perform detailed projections, but make assumptions that are unrealistic when operating in a complex and dynamic world. Ideally, a planning agent in such a world should be able to do both. In order to do this, the agent has to be able to differentiate between those situations in which detailed information would aid it in making its decisions, and those in which such information would not materially improve its performance. We propose an approach to this problem, using well-characterized heuristics to decide what information would be useful, whether to gather it and if so, how.

Introduction

There have been two main approaches to the problem of planning in complex worlds: ignoring the complexity of the world, and ignoring the complexity of the reasoning required to operate in it. So-called "classical" planners are designed to operate in very simple worlds of which they have certain knowledge. (Sussman 1975, Sacerdoti 1977, Chapman 1987). Planning is assumed to occur in advance of execution, and to be subject to no limitations on the time allowed to produce a plan. When there are several possible courses of action, such a planner can thus predict the results of each and choose the one with the best pay-off. These planners can produce plans that are optimal, given the assumptions under which they are working; however, those assumptions are highly unrealistic. So-called "reactive" planners are designed to operate in more complex worlds (Brooks 1986, Agre & Chapman 1987, Firby 1989). Rather than planning in advance, such systems decide *at execution time* on the actions to be undertaken, based solely on the current circumstances. Of course, such systems are not immune from the problematic interactions that projection is designed to discover; but it is assumed that the system's action rules have been written in a way that largely avoids such interactions. The issue of *how* these rules are written has largely gone unaddressed.

* This research was supported in part by DARPA, monitored by AFOSR under contract F49620-88-C-0058, and by AFOSR under contract AFOSR-89-0493. The Institute for the Learning Sciences was established in 1989 with the support of Andersen Consulting, part of The Arthur Andersen Worldwide Organization. The Institute receives additional support from Ameritech, an Institute Partner, and from IBM.

Ideally we would like a planning agent to be able to move quickly enough to take advantage of opportunities as they arise in the world, while at the same time deciding what to do based on projection of the expected consequences of proposed action sequences. If we view the two approaches to planning described above—complete projection on the one hand, pure reactivity on the other—as the ends of a spectrum of possible approaches, we would thus like a planning agent to fall somewhere in the middle, and to be able to adjust its position in the spectrum according to the demands of the moment. Such an agent would thus be required to balance the benefits of projection, in the form of making better decisions, against the costs, in the form of the time and information required. A trade-off must be made between increased accuracy in projection, and the increased cost of planning. In this paper we discuss an approach to this process, concentrating especially on the problem deciding under what circumstances the agent should try to gather more information.

Rock delivery

To see what issues are involved in performing plan projection on the fly, let's consider the robot delivery truck domain used by Firby (1989) and Hanks (1990). Suppose that the truck is parked next to a red rock and a green rock, and has a goal to deliver a red rock to the red rock store, and another goal to deliver a green rock to the green rock store. Should it pick up both rocks now, or should it pick up only one? There are opportunities to further both delivery goals, and it must be decided whether to take advantage of either or both of them. If, for example, the truck had another goal, involving, say, blue rocks, that could not be achieved if the time were spent on picking up red or green rocks, it might be a good decision to pick up neither rock. If the truck were running out of fuel, picking up both rocks (which uses little fuel) then going off to refill its tanks would make sense.

When a planning agent has recognized an opportunity to achieve a goal, it must consider whether taking advantage of the opportunity would be worthwhile. If taking advantage of the opportunity would be relatively costly compared to other possible methods of pursuing the goal whose achievement is facilitated by the opportunity, it may be best to ignore it. On the other hand, if there is an opportunity to do something very cheaply that would provide a large pay-off, this opportunity may be worth pursuing even at the cost of suspending other ongoing plans.

Utility

Traditionally, a goal is a logical expression describing two regions of the outcome space: that in which the goal is satisfied (the set of states with high utility) and that in which it is not (the set of states with low utility). Haddawy & Hanks (1990) extend this notion by viewing a goal as describing a partition of the outcome space such that within each region of the partition, all utility values fall within given bounds. Using this view, we can express the idea that many outcomes, while failing to achieve the stated goal completely, nonetheless have a high utility with respect to that goal. The more complex characterization allows a planning agent to consider the effects of different ways of failing to achieve a goal, as well as providing a simple way of expressing goals that may not have absolute criteria for achievement (such as a goal to deliver as many rocks as possible), and weighing them against other goals in a rational manner.

Although the decision procedure can be summarized by saying "choose the course of action with the highest expected utility" (Feldman & Sproull 1977), the calculation of the expected utilities is often no simple matter, as the following examples demonstrate.

Deadlines

If our delivery planner has to decide whether to pick up the red rock, the green rock, or both, it should compare the utility values arising from the expected delivery times for both rocks in all three cases, and choose the action with the highest associated value. The notion of a deadline can be expressed by saying that all states in the outcome space which involve delivering the rock before a certain time have a high utility associated with them.

Given the expected utility values for the various courses of action, choosing the one that results in the highest expected utility value is a simple optimization problem. Unfortunately, the specific utility functions for, say, the delivery of the red and green rocks will not usually be explicitly available. In order to calculate the expected utility of a particular course of action, we must calculate the probability of each possible outcome associated with that course of action, and the associated utility. Such calculations, which are based on Bayesian theories (Hanks 1990), depend in turn on many values of both prior and conditional probabilities.

In order to decide what to do, a planner must formulate a set of possible courses of action, then project the consequences of each. Each such projection requires information about the possible results of all the component actions. The real problem in choosing a plan is therefore one of information. If sufficient information is available, the decision-making process becomes simple.

Resources

Many of the conditional probabilities involved in evaluating a proposed action sequence are based on reasoning about the resources that will be consumed in performing the sequence. For example, if the red and green rocks are to be delivered to different locations, the relative positions of those locations

can influence the decision as to which rocks to pick up now. If one destination is directly on the way to the other, or requires only a slight detour, picking up both rocks at once is likely to provide significant savings in time, fuel consumption, and wear and tear on the truck over delivering one rock before returning to pick up the other. If, on the other hand, the two locations are far apart and a good route between them passes through the place at which the decision is being made, it may be better to deliver one and then return for the other.

In such cases there are many factors to be balanced, all involving a detailed domain theory. The knowledge that is needed includes facts such as:

- Higher speeds use more fuel, but may help meet a deadline.
- A bad hilly road may be shorter than a good flat road.
- Speeds are lower on hilly roads than on flat roads.
- There is more wear and tear on the truck on bad roads, and fragile contents may be damaged.

Again, we can see that the more information, in the form of a detailed domain theory, that the agent has, the better its decisions can potentially be.

Alternative achievement

When deciding whether to take advantage of an opportunity to achieve a goal, the agent must also consider other available opportunities to achieve that goal, and whether they would have higher overall utility. For example, if the planning agent knows that there is a red rock at the green rock store, an effective course of action might be to deliver the green rock, then pick up a red rock at the time of delivery, rather than pick up a red rock now and carry it around.

The factors that should be taken into consideration in cases like these include:

- Likelihoods of other opportunities existing as expected, and unexpected opportunities arising.
- Costs of taking advantage of the various opportunities.
- Creation of opportunities through substitution: if the top level goal is to fill a container with red objects, for example, the extra benefit obtained by delivering a (large) red rock instead of a (small) red apple might not be worth the extra costs.

In such cases the type of information that helps in the decision-making process is information about the state of the world, including information about the probabilities of various facts being true.

Information and decisions

Although the specification of factors relevant in decision making in the examples given above is not nearly complete, in most cases we would not need to consider the decision in even this much detail. Typically we would believe, for example, that the extra load on the truck from a single rock would affect its performance very little, that the time taken to pick up one rock is unlikely to affect our ability to deliver the other rock on time, and that there are few green rocks in the world so we should take advantage of the opportunity to acquire this one. If these beliefs are true, performing more

detailed reasoning will not improve the decision. In many cases, then, it appears that a decision can be made without considering many of the details of the situation.

However, in cases where decisions are more finely balanced, performing more detailed reasoning could make all the difference. Detailed estimates of timings, routes to be taken and so on, are needed when the deadlines are tight, for example. In order to decide on a course of action and be *sure* that it is the best one available, all possible courses of action must be considered and their possible outcomes projected. In an uncertain world, this projection cannot be performed deterministically, but must use probabilistic reasoning. The projection process needs information about probabilities for a wide range of possible events, about the actual state of the world as observed by the planning agent, and about the utilities of all the actions and outcomes to be considered. The projection process also requires time.

An agent operating in a complex, dynamic world frequently has to make decisions, and often crucial information is unavailable and time is severely limited. There are therefore many occasions when the agent simply cannot perform the detailed projections needed in order to make the best decision. There may well be occasions on which the benefit gained from making the right decision is less than the cost of making the decision. In such cases, there is no net benefit derived from the effort that is put in to making the decision, and overall a better result would be achieved by simply picking a course of action at random.

Gathering information

A planning agent can choose actions in the service of the goal to decide which action to perform next as well as those in direct service of its explicit goals. Among these actions are methods of gaining more useful information about the situation. For instance, suppose the agent has decided to deliver the red rock and is now trying to decide whether to pick up the green one. If there were a telephone available at the destination of the red rock, it might be worth calling ahead to find out whether there is a green rock there, too. Using the telephone presents an alternative to cogitating about the possible presence or absence of green rocks based on the agent's prior knowledge. Of course, the planning agent probably does not want to call all possible places in the world to see if there are green rocks available, but should confine its calls to those places which fit in with its other plans. It needs to have a characterization of the type of information that will be useful in making the decision, and the likelihood of gaining it through performing various actions, so that it can decide which, if any, to perform.

Planners and information gathering

Purely reactive planners choose actions based on features of their environment, with no explicit projection. In effect, they have a set of rules that guess what the result of the projection would be if it were done, in that the rules are intended to make the agent react appropriately. Firby's RAPs (reactive action package) system (1989), for example, has no explicit knowledge of any probabilities that might affect its

behavior. Instead of knowing that an action is especially unreliable, the RAPs system simply doesn't try one that has failed often in the recent past. Agre & Chapman's Pengi (1987) also performs no explicit projection. The system has no preconceived notion of the results of performing an action, but has rules that arbitrate between possible actions. The result of these rules is intended to be that the most appropriate action is chosen, but no representation or search of future world states is performed when applying them. Instead the system's designer has built the results of such reasoning into the rules. Classical planners, on the other hand, explicitly project the results of all the actions that they consider. However, it is assumed that all the required information is available to the system at no cost. Gathering information is thus not an issue for either of these types of systems.

Deciding to gather information

We need some way of looking at the utility obtained from gathering information in terms of the role played by information in improving decisions. We also need to be able to estimate the costs involved in gathering it. Horvitz *et al* (1989), have looked at the possibility of controlling decision-theoretic inference by considering the expected value of such inference in relation to the decision to be made. Their method uses the properties of *belief networks* (Howard & Matheson 1981, Pearl 1986), and in particular the properties of the algorithm they use to evaluate the networks.

Unfortunately, reliance on a particular algorithm for evaluating belief networks is unlikely to be very useful in more dynamic domains, where the agent can find itself in totally new situations. As we discussed earlier, there is a large amount of information required in order to make good projections of the expected utility of the possible outcomes. If belief networks are used to perform the projection, a new network would have to be constructed for each decision in order to take the relevant influences into account. Constructing the necessary dependencies would be prohibitively expensive, especially as the decisions must be made at execution time.

An alternative approach is suggested by viewing information gathering as just one of the many possible tasks that can be performed at any particular time, and looking at how existing planners decide between tasks. Many systems use *heuristics* to make decisions when they cannot perform a complete analysis. Work has also been done on approximate analysis (Hanks 1990, Wellman 1990). The success of these systems suggests that heuristics and approximations can provide useful information that can help in making a decision. If we view information gathering as a planning task, it is reasonable to ask whether the sorts of decision heuristics employed by planners can also be applied to the problem of deciding when to gather more information. We believe that the answer is yes.

However, in general the heuristics used by existing planners are at best *ad hoc*, and are often implicit in the control structure. We intend to characterize useful heuristics and approximations in a more principled way, so that we can understand when and how they should be used. This characterization will have to specify the results produced by the

heuristics in terms of the quality of the decisions that use them, and the costs of using them in terms of time and other resources. We will investigate their use in the context of planning in a complex and dynamic domain so that their applicability can be demonstrated.

Effective independence

When a planner is considering whether to perform an action, it should consider the benefits and costs associated with the action. The benefits of the action can be measured through its (positive) effect on the goal that it directly furthers and on other goals. The costs can be measured through its (negative) effects on its own and other goals. The resources that are consumed by performing the action affect goals inasmuch as they cannot then be used by other actions. The calculation of these benefits and costs requires information, as we have discussed above, and we intend to investigate methods of judging what information would be useful, how that information can be gathered, and whether it is worth gathering. We have not yet considered the question of what effect an action has on the goal that it directly furthers, and the information required in order to judge that effect, but have started by looking at an action's effects on other goals.

Our first step is to come up with heuristics that indicate when the action's effects on goals that it does not directly further can be ignored. In this case the planner does not need information enabling it to calculate those components of the utility of the action. On the other hand, if the effects on other goals are significant, the information relating to them is required in order to make a good decision.

Clearly, if an action is independent of all goals that it does not directly further, the planner can indeed ignore its effects on other goals. Most actions are not, unfortunately, independent of all other goals. In our delivery example, for instance, the goals to deliver the red and green rocks interact in important ways. These interactions are responsible for most of the complications we discuss. However, in many cases such interactions are insignificant. When we say that the exact time taken to pick up the green rock is unimportant because it is very short compared to the time available for the delivery of the red rock, and that the weight of the green rock can be ignored as being negligible compared to the overall weight of the truck, we are saying that we can treat the action of picking the green rock up as though it were independent of the goal to deliver the red rock. In this case we say that the action of picking the green rock up is *effectively independent* of the goal to deliver the red rock. When actions are not effectively independent, we need to reason more about the interactions. Heuristics to decide when actions are effectively independent would therefore be very useful in determining whether we need more information about the situation.

Violations of effective independence

At present we are considering two principal ways in which independence may be violated: through preconditions and through post-conditions.

Two of the more common types of preconditions are that the agent has sufficient resources to carry out a plan that

would achieve the goal, and that certain states of the world are true before a task is performed. In many cases, the resources needed in order to achieve a goal are plentiful. However, conflicts can arise when they are limited:

- If the truck were nearly full, picking up the green rock would leave no room for anything else. This has no effect on the goal to deliver the red rock, which is already in the truck, but would mean that other delivery tasks could not be undertaken.
- If picking up a rock uses a large amount of fuel, and the truck's tank is nearly empty, there may not be enough to pick up the green rock and deliver the red rock without refilling the tank. Thus the goal to deliver the red rock could be affected in a couple of ways: the truck could run out of fuel, so preventing the delivery, or could refuel along the way, thus delaying the delivery.

Another type of independence violation through preconditions arises from the fact that performing an action can sometimes change the expected state of the world in a way that interferes with the achievement of another goal. Elements of the world state that can be affected include the weight of the truck, time, and location:

- Suppose that the truck's being under a certain weight is a precondition of being able to cross a bridge which is on the route to deliver the red rock. If the green rock is very heavy, picking up the green rock would violate that precondition by changing the state of the truck so that the bridge would no longer bear it.
- If there is a deadline for the delivery of the red rock, a precondition for the delivery is that the transporting task should not start too late. Picking up the green rock would delay the time at which the truck could start on the task of taking the red rock to its destination, possibly violating the precondition if the deadline is tight.
- Even if there is no green rock at the same place as the red rock, there is always an opportunity to move elsewhere to collect a green rock before continuing with the task of delivering the red rock. The change of location would normally significantly affect the task of transporting the red rock, which would have to start from a place other than that which was expected.

An action can affect another goal by changing the state of the world by making it impossible to achieve the goal:

- Suppose that there is only one green rock, but that there are also some green apples. The truck has goals to deliver a green rock and a green object to separate locations. If it tries to achieve the second goal by delivering the green rock, it has made the first goal unachievable.

Of course, in the last example, it is still possible to achieve the goal of delivering a green rock. In order to do this, however, the goal of delivering a green object would become unsatisfied: a post-condition of the goal is that the green object should be left at the location to which it is delivered, and not moved away from that place. In this situation the agent would need to acquire information about the ways in which both goals can be satisfied, by using a green apple to achieve the goal to deliver a green object, for example.

Detecting threats to independence

There are three main ways in which the threats to independence that we have discussed can be detected:

- An action may be known to be liable to have a certain effect on goals. For example, moving to another location will always cause a change of state, and should prompt the agent to look for possible conflicts.
- A goal may be known to be sensitive to certain types of interference. If the red rock is particularly fragile, for instance, the agent should beware of actions involving objects that might damage it. The principle involved is very similar to the notion of protections used in classical planning (Sussman 1975, Sacerdoti 1977). Instead of a violation being fatal, however, in this case it simply means that there may be a problem and that the matter should be investigated further.
- A resource or state may be prone to conflicts. For example, if there is little fuel, or any other resource is in short supply, the agent should be on the look out for actions that affect that resource.

A planning agent can use these strategies in a two-stage process to determine if there is effective independence. The first stage uses heuristics to decide whether to perform the second, more detailed, stage.

In order to judge whether an action and a goal are effectively independent, the planning agent should consider aspects of the action, goal, and world state. The agent should perform more detailed analysis in the following cases:

- Action-driven
 - There is a resource which the action uses heavily.
 - There is a world state which the action changes significantly.
- Goal-driven
 - The achievement of the goal relies on a resource used by the action.
 - There are any world-states that are fatal to the achievement of the goal.
- World-driven
 - There are any resources that are in short supply.

Conclusion

An intelligent agent must be able to differentiate between important decisions and those that have few or unimportant consequences. The complexities involved in such reasoning in a dynamic world, in which exactly the same situation will probably never arise twice, are enormous. A decision-maker therefore must be able to reason about the type of decision required in a particular set of circumstances and the information that is needed in order to be able to make that decision effectively.

This type of decision behavior depends on the ability to estimate how good we expect a decision to be, and the expected benefits of making it better. This means that we need ways of characterizing the quality of decisions, how the information used in them affects that quality, and the availability of relevant information. All these estimates must in turn be related to the expected cost of gathering the information.

An agent reasoning in this way about decision making must be able to do so efficiently, otherwise the costs of the reasoning will outweigh the benefits.

As we discussed above, the reasoning involved in making decisions about planning can be very complex and time-consuming, and rely on large amounts of information which may not be readily available. One approach to carrying out this reasoning in a timely manner is to develop heuristics and approximations that can be used when detailed reasoning is inappropriate. Such heuristics must also be used when considering whether it is worth gathering more information, either by performing more reasoning or by other methods or whether the less exact methods will be more appropriate in the circumstances.

We intend to continue our investigation of effective independence as one approach to the problem of characterizing useful information in the context of planning decisions. Effective independence concerns the relationship between an action and a goal it does not directly further, and we will look at other aspects of that relationship as well as the way in which an action affects the goal that it serves.

Acknowledgements: Thanks to Larry Birnbaum, Matt Brand, Mike Freed and Bruce Krulwich for many useful discussions.

References

- Agre, P. E., & Chapman, D. (1987). Pengi: An implementation of a theory of activity. Proceedings of AAAI-87.
- Brooks, R. (1986). A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, vol 2, no 1.
- Chapman, D. (1987). Planning for Conjunctive Goals. *Artificial Intelligence*, 32, 333-337.
- Feldman, J. A., & Sproull, R. F. (1977). Decision Theory and Artificial Intelligence II: The Hungry Monkey. *Cognitive Science* 1, 158-192.
- Firby, R. J. (1989). Adaptive execution in complex dynamic worlds. Technical Report YALEU/CSD/RR #672, Dept. of Computer Science, Yale University, New Haven CT.
- Haddawy, P., & Hanks, S. (1990). Issues in Decision-Theoretic Planning: Symbolic Goals and Numeric Utilities. Innovative Approaches to Planning, Scheduling and Control: Proceedings of a Workshop. DARPA, San Diego.
- Hanks, S. J. (1990). Projecting plans for uncertain worlds. Technical Report YALEU/CSD/RR #756, Dept. of Computer Science, Yale University, New Haven CT.
- Horvitz, E. J., Cooper, G. F., & Heckerman, D. E. (1989). Reflection and Action Under Scarce Resources: Theoretical Principles and Empirical Study. Proceedings of the Eleventh International Joint Conference on Artificial Intelligence, Detroit, MI.
- Howard, R. A., & Matheson, J. E. (1981). Influence Diagrams. In R. A. Howard & J. E. Matheson, editors, *Readings on the Principles and Applications of Decision Analysis*, vol II, pp721-762. Strategic Decisions Group, Menlo Park, CA.
- Pearl, J. (1986). Fusion, Propagation, and Structuring in Belief Networks. *Artificial Intelligence* 29, 241-288.
- Sacerdoti, E. (1977). *A Structure for Plans and Behavior*. American Elsevier, New York.
- Sussman, G. J. (1975). *A Computer Model of Skill Acquisition*. American Elsevier, New York.
- Wellman, M. P. (1990). *Formulation of Tradeoffs in Planning under Uncertainty*. Pitman, London and Morgan Kaufman, San Mateo, CA.