An Investigation of Probabilistic Interpretations of Heuristics in Plan Recognition

Peter van Beek
Department of Computing Science
University of Alberta
Edmonton, Alberta, Canada T6G 2H1
vanbeek@cs.ualberta.ca

Abstract

Plan recognition is the process of inferring a plausible set of plans that explain an agent's actions. In this paper, I use a small example to focus on some non-probabilistic heuristics proposed in the literature for preferring one plan over another. I show some of the conditions or constraints on the probability distributions so that the plan that is preferred by the heuristics is also the most probable plan and I look at some of the implications of these conditions. I also show that if the conditions do not hold there exist cases where the results of the heuristics clash with that of probabilities. One of the most interesting results of the analysis is that, given the assumption that the plan library is complete, the heuristics examined can be given a probabilistic interpretation or justification if and only if any two basic plans in the plan library which share a step are equally likely. The usefulness of the analysis is that we can test whether the conditions hold in a particular domain and so gain more insight into whether our choice of heuristic is appropriate to that domain. Further, this work can be seen as providing an alternative justification of the heuristics.

Introduction

Knowing the plan an agent is pursuing is useful if we wish to intentionally help or hinder that agent. For example, in question-answering systems recognizing the plan underlying a user's query aids in both understanding the query (e.g., (Carberry 1988)) and in generating an appropriate response (e.g., (Allen 1983; Cohen, Schmidt, & van Beek 1994; McKeown, Wish, & Matthews 1985)). Assuming, however, that we have no direct access to an agent's plan, the best that we can do is to postulate one or more plans, based on observation of the agent's actions, that are plausible candidates for being the agent's plan. The process of inferring a plausible set of plans from observation of an agent's physical or speech actions is called plan recognition.

A plan is said to explain a set of actions if it contains those actions. Given a representation of the possible plans in a domain and a set of observed or described actions, the plan recognition problem in its most general form is to determine the set of all plans that explain the actions. Four principal methods for plan recognition have been proposed in the literature. The methods are plausible inference (Allen 1983; Calistri 1990; Carberry 1988; 1990; Charniak & Goldman 1993; Litman & Allen 1987), parsing (Sidner 1985; Vilain 1990), circumscribing a hierarchical representation of plans and using deduction (Kautz 1987), and abduction (Charniak & McDermott 1985; Konolige & Pollack 1989; Lin & Goebel 1991; Poole 1989). The set of all plans that explain a set of actions may be large and many methods for selecting or preferring one plan over another have been proposed in the literature: action- and search-based heuristics (Allen 1983), a focusing heuristic based on discourse coherence considerations (Carberry 1985), levels of likelihood (McKeown, Wish, & Matthews 1985), probabilities (Bauer *et al.* 1993; Carberry 1990; Charniak & Goldman 1993; Charniak & McDermott 1985; Neufeld 1989), asking the user to distinguish (Cohen, Schmidt, & van Beek 1994), temporal consistency (Kautz 1987), assuming as few top level goals as is consistent (Carver, Lesser, & McCue 1984; Kautz 1987), and abductive or theory preference heuristics such as specificity (Konolige & Pollack 1989; Lin & Goebel 1991; Poole 1985).

In this paper, I focus on non-probabilistic heuristics proposed by Carberry (1985), Carver et al. (1984), McKeown et al. (1985), and Kautz (1987) for preferring one plan over another. These four heuristics, while varying in motivation, broadly agree on which plans to prefer. I take as my starting point the view that probability theory is a normative theory of plausible reasoning. As such, I use it to analyze the four heuristics. I examine, using a small example, some conditions under which the heuristics can be given a probabilistic interpretation; that is, I show some of the conditions or constraints on the probability distributions so that the plan that is preferred by the heuristics is also the most probable plan. I also show that if the conditions do not hold there exist cases where the results of the heuristics clash with that of probabilities. One of the most interesting results of the analysis is that, given the assumption that the plan library is complete, the heuristics examined can be given a probabilistic interpretation or justification if and only if any two basic plans in the plan library which share a step are equally likely. The usefulness of the analysis is that we can test whether the conditions hold in a particular domain and so gain more insight into whether our choice of heuristic is appropriate to that domain. This is important as the heuristics are applied every time we get a new observation of an agent's actions, and the output of the plan recognition systems—the plan that the systems posit as the plan that the user is pursuing—depends on the heuristic. If we are going to use the output of a plan recognition system to, for example, provide some sort of automatic plan completion, we will want it to be correct as often as possible. As well, this work can be seen as providing an alternative justification of the heuristics.

Four heuristics and a plan library

In this section I briefly describe the heuristics proposed by Carberry (1985), McKeown et al. (1985), Carver et al. (1984), and Kautz (1987). All four of these heuristics are methods for combining multiple observations. I then give a small plan library which is used to illustrate the analysis.

Carberry (1985) gives a focusing heuristic that is motivated by discourse coherence considerations. For each new utterance (observation) a set of candidate plans is hypothesized. The focusing heuristic is then used to select a "best" candidate and incorporate it into a structure, called a context model, that represents the plan inferred from the preceding dialogue.

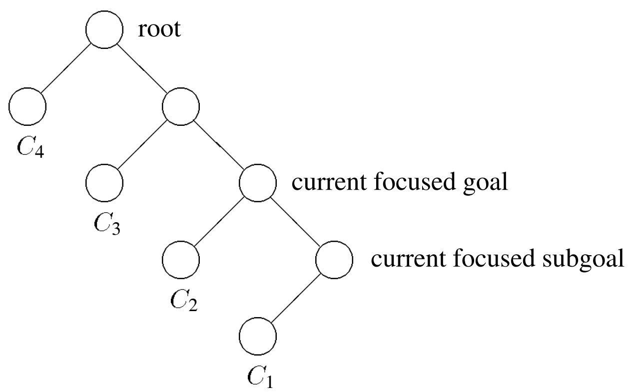

Figure 1: Illustration of Carberry's heuristic

A focused goal and a focused subgoal—the most recently expanded subgoal of the focused goal—are tracked. The heuristic is to test first if a candidate is part of an expansion of a plan for the current focused subgoal and, if not, to next try the current focused goal, on up to the root, stopping after the first successful test. So, in the example above, the candidate plans would be tested at position C1, then at C2, on up until a test succeeds. If all of these tests fail, we test whether the root is an expansion of a candidate plan or if some other plan can be expanded to include both a candidate and the root. It should be noted that in my analysis I do not capture these preferences—I simplify the analysis by assuming all of the C_i are equally preferred.

McKeown et al. (1985) give a similar heuristic. They describe their heuristic as essentially determining the lowest common ancestor of the candidate plans determined from the new observation and the plan determined from the previous observations.

Carver et al. (1984) propose a heuristic that prefers to fit a new observation into a plan already in progress over initiating a new top-level goal.

Kautz (1987) gives a similar heuristic that prefers plans which require the agent to have as few unrelated intentions as possible. Without going into the details of the representation for plans that Kautz develops, the heuristic can also be expressed as preferring the plan with the smallest number of end events, where an end event is an event that is pursued for itself.

I use a small library of possible plans to investigate some of the underlying conditions such that the four heuristics can be given a probabilistic interpretation. The example library is shown in Figure 2. In this example, there are three end events (d_1, d_2, and d_3) and two actions (s_1 and s_2). The arrows represent decomposition. Each plan has exactly one decomposition so there are also only three basic plans. For example, the plan to accomplish end event d_2 consists of both action s_1 and action s_2. I refer to a basic plan by the name of its end event. I also allow an agent's plan to be a conjunction of basic plans such as, for example, the plan $d_1 \wedge d_2$.

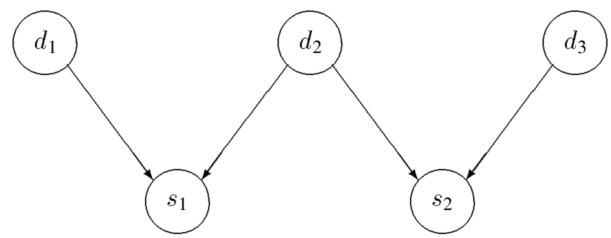

Figure 2: Small plan library

Examining the heuristics

In this section I examine all of the cases within the small plan library where the heuristics prefer one plan over another. Each of the cases contributes new constraints to the set of overall constraints which must be satisfied for the heuristics to be given a probabilistic interpretation. As a corollary, I also demonstrate that there are reasonable cases where the heuristics do not favor the most probable plan. Three of the cases are examined in detail. To make the intuition more clear, consider the interpretation of the plan library shown in Figure 3.

In Case 1, one plan explains all of the observations and the other plan does not. Here the heuristics favor the plan that does explain all of the observations. As an example, suppose we observe an agent getting a gun and going to the bank (see Figure 3). It seems reasonable to assume the agent is robbing the bank rather than going hunting.

In Case 2, both plans explain all of the observations but one plan is larger (in terms of the number of end events) than the other. Here the heuristics favor the smaller plan. As an example, again suppose we observe an agent getting a gun and going to the bank. Without any knowledge of the agent performing the action or any other background knowledge, it seems reasonable to assume that the agent is robbing the bank. We prefer the simpler explanation for the agent's actions rather than the more complex explanation where the agent is going hunting and going to the bank to cash a check. Preferring the simpler explanation is consistent with Occam's razor—a well-known scientific and philosophical rule which states that entities should not be multiplied unnecessarily and so we should prefer the simpler explanation over the more complex explanation. However, suppose that we are supplied with the background knowledge that the agent is a law enforcement officer. It now seems that the more complex explanation may be preferred. Contrast this with the background knowledge that the agent is an anti-vivisectionist which seems to reinforce the belief that robbing a bank is the "right" answer. I examine this in more detail below.

In Case 3, both plans explain all of the observations and both plans are of equal size (in terms of the number of end events). Here the heuristics do not distinguish between the two plans; that is, both plans are equally preferred. As an example, suppose we observe an agent getting a gun. It seems reasonable to not prefer the plan where the agent is going hunting over the plan where the agent is robbing the bank, and vice versa, and to perhaps wait for further observations to help distinguish between the two.

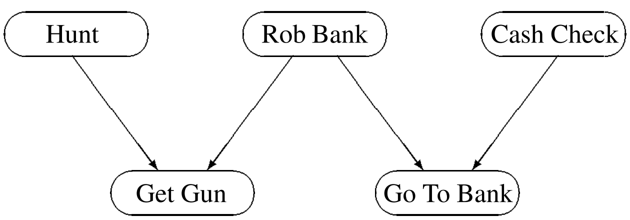

Figure 3: One interpretation of the plan library (adapted from (Kautz 1987))

Before we examine the heuristics in detail, we first need to specify a probabilistic model. A probabilistic model is a specification of probabilistic information such that the probability of every well-formed sentence can be computed. Our example is small enough that we can specify a probabilistic model by specifying a joint probability distribution. That is, we assign a probability to every elementary event in the language such that the probabilities p_i sum to one.

$$
\begin{aligned}
&P(\neg d_1, \neg d_2, \neg d_3, \neg s_1, \neg s_2) = p_0 \\
&P(\neg d_1, \neg d_2, \neg d_3, \neg s_1, s_2) = p_1 \\
&P(\neg d_1, \neg d_2, \neg d_3, s_1, \neg s_2) = p_2 \\
&\ldots \\
&P(d_1, d_2, d_3, s_1, \neg s_2) = p_{30} \\
&P(d_1, d_2, d_3, s_1, s_2) = p_{31}
\end{aligned}
$$

Such a set of mutually exclusive and exhaustive events is often called a frame of discernment.

Almost all work on plan recognition makes the assumption that the plan library is complete in the sense that it is assumed to contain all valid events and all valid relationships between events. Pollack (1986) was the first to recognize that this assumption was implicit in previous work and to characterize it as a type of closed-world assumption. While recognizing that it can be an overly strong assumption in some domains, Pollack also makes this assumption in her work. Kautz (1987) also explicitly makes this assumption and an important contribution of his work is a formalization of this sort of closed-world assumption. I will make the probabilistic equivalent of Kautz's assumptions. I show how making these assumptions strongly constrains the joint probability function. I wish to be clear about why I am making the probabilistic equivalent of the closed-world assumption. The aim of this work is to determine when the four non-probabilistic heuristics can be given a probabilistic interpretation *as they are used in existing plan recognition systems*. The four heuristics which I am examining in this paper were all proposed in the context of a system or a framework and in all four of these systems the closed-world assumption is made either explicitly or implicitly. Note that at the moment we make no additional assumptions such as unconditional or conditional independence or mutually exclusive top-level events.

Assumption 1: The plan library is assumed to contain all possible ways to achieve an end event. That is, every possible decomposition of a plan is represented in the plan library. In our example plan library, this is equivalent to asserting that $P(s_1 \mid d_1) = 1$, $P(s_1 \wedge s_2 \mid d_2) = 1$, and $P(s_2 \mid d_3) = 1$. In words, the assertion $P(s_1 \mid d_1) = 1$ says that, given that we are achieving event d_1 it is a certainty that we are also performing action s_1, since that is the only way to accomplish d_1. Together these assertions constrain many of the p_i to be zero. To illustrate,

$$
\begin{aligned}
&P(s_2 \mid d_3) \\
&\quad = \frac{P(s_2 \wedge d_3)}{P(d_3)} \\
&\quad = \frac{\alpha}{\alpha + p_4 + p_6 + p_{12} + p_{14} + p_{20} + p_{22} + p_{28} + p_{30}} \\
&\quad = 1,
\end{aligned}
$$

where $\alpha = p_5 + p_7 + p_{13} + p_{15} + p_{21} + p_{23} + p_{29} + p_{31}$ and where in the first step we used the definition of conditional probabilities and in the second the specified joint probability distribution. Hence,

$$p_4, p_6, p_{12}, p_{14}, p_{20}, p_{22}, p_{28}, p_{30} = 0.$$

To make clear the intuition, let us consider as an example why p_{30} must equal 0 given our assumption; i.e., why $P(d_1, d_2, d_3, s_1, \neg s_2) = p_{30} = 0$. In words, this elementary event is purportedly a plan to, among other things, achieve d_3 but without performing the action s_2. But since s_2 is the only possible way to achieve d_3, this event is impossible.

Assumption 2: The plan library is assumed to contain all possible reasons to perform an action and it is assumed that there are no useless actions. In our example plan library, this is equivalent to asserting that $P(d_1 \vee d_2 \mid s_1) = 1$ and

$P(d_2 \vee d_3 \mid s_2) = 1$. In words, the assertion $P(d_1 \vee d_2 \mid s_1) = 1$ says that, given that we are performing action s_1 it is a certainty that we are also achieving either d_1 or d_2 or both, since those are the only reasons to perform s_1. These assertions also constrain many of the p_i to be zero. To illustrate,

$$
\begin{aligned}
& P(d_1 \vee d_2 \mid s_1) \\
& \quad = \frac{P((d_1 \wedge s_1) \vee (d_2 \wedge s_1))}{P(s_1)} \\
& \quad = \frac{\alpha}{\alpha + p_2 + p_3 + p_7} \\
& \quad = 1,
\end{aligned}
$$

where $\alpha = p_{11} + p_{15} + p_{18} + p_{19} + p_{23} + p_{27} + p_{31}$. Hence,

$$p_2, p_3, p_7 = 0.$$

To make clear the intuition, let us consider as an example why p_2 must equal 0 given our assumption; i.e., why $P(\neg d_1, \neg d_2, \neg d_3, \neg s_1, s_2) = p_2 = 0$. In words, this elementary event says that action s_2 is being performed but that it does not accomplish any end event. But since no action is useless, this event is impossible.

Given the above closed-world assumptions, the following elementary events have non-zero probabilities (where I have renamed the probabilities on the right hand sides).

$$
\begin{aligned}
& P(\neg d_1, \neg d_2, \neg d_3, \neg s_1, \neg s_2) = p_0 \\
& P(\neg d_1, \neg d_2, d_3, \neg s_1, s_2) = p_1 \\
& P(\neg d_1, d_2, \neg d_3, s_1, s_2) = p_2 \\
& P(\neg d_1, d_2, d_3, s_1, s_2) = p_3 \\
& P(d_1, \neg d_2, \neg d_3, s_1, \neg s_2) = p_4 \\
& P(d_1, \neg d_2, d_3, s_1, s_2) = p_5 \\
& P(d_1, d_2, \neg d_3, s_1, s_2) = p_6 \\
& P(d_1, d_2, d_3, s_1, s_2) = p_7
\end{aligned}
$$

Let us now examine in detail the three cases outlined above.

Case 1: If one plan explains all of the observations and another plan does not explain all of the observations, the heuristics favor the plan that does explain all of the observations. Thus, in order for there to be a probabilistic interpretation of the heuristics, we must consistently constrain the joint probability distribution such that the probability of a plan that does explain all of the observations is greater than the probability of a plan that does not. For example, suppose that our plan library is as shown in Figure 2 and that we observe both s_1 and s_2. The heuristics all prefer plan d_2 over the plan d_1. Thus,

$$
\begin{aligned}
& P(d_2 \mid s_1 \wedge s_2) \\
& \quad = \frac{P(d_2 \wedge s_1 \wedge s_2)}{P(s_1 \wedge s_2)} \\
& \quad = \frac{p_2 + p_3 + p_6 + p_7}{p_2 + p_3 + p_5 + p_6 + p_7}
\end{aligned}
$$

and

$$
\begin{aligned}
& P(d_1 \mid s_1 \wedge s_2) \\
& \quad = \frac{P(d_1 \wedge s_1 \wedge s_2)}{P(s_1 \wedge s_2)} \\
& \quad = \frac{p_5 + p_6 + p_7}{p_2 + p_3 + p_5 + p_6 + p_7}
\end{aligned}
$$

and thus,

$$P(d_2 \mid s_1 \wedge s_2) > P(d_1 \mid s_1 \wedge s_2)$$

if,

$$p_2 + p_3 > p_5.$$

That is, the constraint $p_2 + p_3 > p_5$ must be satisfied for the heuristics to favor the plan with the highest probability [1].

Case 2a: If two plans both explain all of the observations, the heuristics favor the plan with the smallest cardinality. Thus, we must consistently constrain the joint probability distribution such that the probability of the plan with the smaller cardinality is greater than the probability of the plan with the larger cardinality. For example, suppose that our plan library is as shown in Figure 2 and that we observe s_1 and s_2. The heuristics all prefer plan d_2 over the plan $d_1 \wedge d_3$ (the result is independent of the order the observations are made in)[2]. Thus,

$$
\begin{aligned}
& P(d_2 \mid s_1 \wedge s_2) \\
& \quad = \frac{P(d_2 \wedge s_1 \wedge s_2)}{P(s_1 \wedge s_2)} \\
& \quad = \frac{p_2 + p_3 + p_6 + p_7}{p_2 + p_3 + p_5 + p_6 + p_7}
\end{aligned}
$$

and

$$
\begin{aligned}
& P(d_1 \wedge d_3 \mid s_1 \wedge s_2) \\
& \quad = \frac{P(d_1 \wedge d_3 \wedge s_1 \wedge s_2)}{P(s_1 \wedge s_2)} \\
& \quad = \frac{p_5 + p_7}{p_2 + p_3 + p_5 + p_6 + p_7}
\end{aligned}
$$

and thus,

$$P(d_2 \mid s_1 \wedge s_2) > P(d_1 \wedge d_3 \mid s_1 \wedge s_2)$$

if,

$$p_2 + p_3 + p_6 > p_5.$$

[1] In some of the literature, especially the diagnostic literature, we only compute the probability of mutually exclusive and exhaustive hypotheses such as $P(\neg d_1, d_2, \neg d_3 \mid s_1 \wedge s_2)$ and $P(d_1, \neg d_2, \neg d_3 \mid s_1 \wedge s_2)$. I am still not sure whether this is what is wanted in this context. It seems to me that what most people in plan recognition are saying when they give a method for performing plan recognition is that the plan that is proposed by the method is one plausible plan for an agent but not excluding the possibility that the agent is performing other plans.

[2] In the general case, Carberry's and McKeown's heuristics are sensitive to the order of the observations. Kautz represents temporal information in the plan library and allows temporal information to be part of the input. Thus all three can capture that different plans are sometimes suggested depending on the order of the observations. This is not straightforward to capture in a probabilistic setting.

That is, the constraint $p_2 + p_3 + p_6 > p_5$ must be satisfied for the heuristics to favor the plan with the highest probability. This constraint can be stated without reference to the observations. We simply want that $P(d_2) > P(d_1 \wedge d_3)$. In words, it must be more likely that we are pursuing the single plan d_2 than that we are pursuing the conjunctive plan d_1 and d_3. It is clear that this may or may not hold in any particular domain, and thus the heuristics may or may not have a probabilistic interpretation or justification in that domain.

Case 3: If two plans both explain all of the observations and both plans are of equal size (in terms of the number of end events) the heuristics do not distinguish between the two. That is, they are equally preferred explanations. Thus, we must consistently constrain the joint probability distribution in such a case so that the probabilities of the two plans are equal. For example, suppose that our plan library is as shown in Figure 2 and that we observe s_1. The heuristics all equally prefer plan d_1 and plan d_2. Thus,

$$P(d_1 \mid s_1) = P(d_2 \mid s_1)$$

if,

$$p_2 + p_3 = p_4 + p_5.$$

That is, the constraint $p_2 + p_3 = p_4 + p_5$ must be satisfied for the heuristics to equally favor these two plans with equal cardinality that both explain the observation.

The three cases above illustrate how constraints on the probability distribution are derived. The complete set of non-redundant constraints is as follows (constructed by determining for each possible observation, s_1, s_2, and $s_1 \wedge s_2$, the preferred explanation(s) from the set of all possible explanations d_1, d_2, d_3, $d_1 \wedge d_2$, $d_1 \wedge d_3$, $d_2 \wedge d_3$, and $d_1 \wedge d_2 \wedge d_3$, and the associated constraints).

$$\begin{array}{l} p_1 + p_3 > p_6 \\ p_2 + p_3 > p_5 \\ p_2 + p_6 > p_5 \\ p_4 + p_6 > p_3 \\ p_1 + p_5 = p_2 + p_6 \\ p_2 + p_3 = p_4 + p_5 \\ p_0 + p_1 + p_2 + p_3 + p_4 + p_5 + p_6 + p_7 = 1 \end{array}$$

It can be seen that if all of the elementary events are equally likely, all of the constraints are satisfied and the heuristics will favor the plan that is the plan with the highest probability. Thus, one way for the heuristics to be given a probabilistic interpretation is if we can assume that the elementary events are equally likely. This assumption can be seen to follow from a powerful and general statistical principle called the principle of maximum entropy (see (Davis 1990) and references therein). The principal of maximum entropy states that if there are n elementary events $e_1, \ldots, e_n$ which form a frame of discernment, we should assign probabilities to the elementary events in such a way that, subject to any constraints on the probabilities we may know, the values give us the least information about which of the events occurs. In our setting, if we assume that the only information we have is that some of the events are impossible, then the minimum information assumption is to assume that each elementary event that is not impossible is equally likely. Thus, the principle of maximum entropy can be viewed as an alternative justification of the heuristics.

A consequence of the principal of maximum entropy, and in particular a consequence of making each elementary event equally likely, is that the basic plans are independent of each other. Independence of basic plans means that knowing that one plan is occurring does not change the probability that another plan is also occurring. In some domains this may be a reasonable assumption, in others it is not. I discuss this further in the last section.

It is also possible for the elementary events to not be all equally likely and for the constraints still to be satisfied. For example, here is one such assignment of probabilities that satisfies the constraints: $p_0 = 0$, $p_1 = 8/32$, $p_2 = 4/32$, $p_3 = 5/32$, $p_4 = 7/32$, $p_5 = 2/32$, $p_6 = 6/32$, and $p_7 = 0$. However, even though the elementary events are not all equally likely, the basic plans are all equally likely; that is, $P(d_1) = P(d_2) = P(d_3)$. It can be seen that for this particular plan library, there is a probabilistic justification or interpretation of the heuristics if and only if all of the basic plans are equally likely. In particular, if not all basic plans are equally likely, then no probabilistic justification or interpretation can be given to the heuristics and, in fact, the plans preferred by the heuristics clash with those preferred by a probabilistic criteria.

The above analysis was done with respect to a small example plan library. What can be concluded about the heuristics applied to arbitrary plan libraries? Consider the set of all constraints shown above. In general, equality constraints will be asserted between pairs of plans for which the heuristics do not prefer one over the other. Such equality constraints will be asserted between every pair of basic plans that share a step. Thus in general we can conclude that there is a probabilistic justification or interpretation of the heuristics *if and only if* every pair of basic plans which share a step are equally likely. This is a strong condition, for as was shown for the small example plan library, one plan may share a step with a second plan which in turn shares a step with a third plan and as a result all three plans must be equally likely.

The plan library as a Bayesian network

In general, specifying a joint probability distribution can require us to specify many probabilities. A technique for reducing the number required is to explicitly capture actual or assumed irrelevance by representing conditional independence relationships. One way to do this is by using a graphical representation of a joint probability distribution called Bayesian networks. Bayesian networks are directed acyclic graphs in which the nodes represent random variables (in our context, the variables are propositional variables that can be either true or false). An arc between two variables in a Bayesian network means that the one variable directly influences the other variable and the strength of the influence is expressed by conditional probabilities. The overall structure of the network encodes the conditional independence assumptions (see (Charniak 1991; Pearl 1988)).

Bayesian networks have been applied to plan recognition. Neufeld (1989) suggests interpreting the plan library as a Bayesian network. Similarly, Charniak and Goldman (1993) use the observations to retrieve candidate plans and then construct a Bayesian network for plan recognition from the candidate plans. Bayesian updating is then used to choose the most likely interpretation for the set of observed actions. Let us follow this work and interpret our example plan library as a Bayesian network. In our example, this is equivalent to making the following assumption.

Assumption 3: The basic plans in the plan library are assumed to be independent.

Bayesian networks encode certain conditional independence assumptions which make it straightforward to write down a joint probability distribution. By inspection of the network in Figure 2 we can immediately write down that,

$$\begin{aligned} & P(d_1, d_2, d_3, s_1, s_2) \\ &= P(d_1)P(d_2)P(d_3)P(s_1 \mid d_1, d_2)P(s_2 \mid d_2, d_3). \end{aligned}$$

Let us make the same closed-world assumptions as before. Once again many of the probabilities of the elementary events are forced to be zero and the probabilities of the elementary events that are not necessarily zero can be simplified. For example, using the equation above we get that,

$$\begin{aligned} & P(d_1, \neg d_2, \neg d_3, s_1, s_2) \\ &= P(d_1)P(\neg d_2)P(\neg d_3) \\ &\qquad P(s_1 \mid d_1, \neg d_2)P(s_2 \mid \neg d_2, \neg d_3) \\ &= 0 \end{aligned}$$

since $P(s_2 \mid \neg d_2, \neg d_3) = 0$ by Assumption 2. The complete results are shown below.

$$\begin{aligned} & P(\neg d_1, \neg d_2, \neg d_3, \neg s_1, \neg s_2) = P(\neg d_1)P(\neg d_2)P(\neg d_3) \\ & P(\neg d_1, \neg d_2, d_3, \neg s_1, s_2) = P(\neg d_1)P(\neg d_2)P(d_3) \\ & P(\neg d_1, d_2, \neg d_3, s_1, s_2) = P(\neg d_1)P(d_2)P(\neg d_3) \\ & P(\neg d_1, d_2, d_3, s_1, s_2) = P(\neg d_1)P(d_2)P(d_3) \\ & P(d_1, \neg d_2, \neg d_3, s_1, \neg s_2) = P(d_1)P(\neg d_2)P(\neg d_3) \\ & P(d_1, \neg d_2, d_3, s_1, s_2) = P(d_1)P(\neg d_2)P(d_3) \\ & P(d_1, d_2, \neg d_3, s_1, s_2) = P(d_1)P(d_2)P(\neg d_3) \\ & P(d_1, d_2, d_3, s_1, s_2) = P(d_1)P(d_2)P(d_3) \end{aligned}$$

So, we need to specify only the priors of the three basic plans in order to be able to compute the probability of any well-formed sentence.

Let us now examine a case similar to that of Case 2a except that here we have two mutually exclusive hypotheses with different cardinalities (based on counting the variables that are true).

Case 2b: Suppose that our plan library is as shown in Figure 2 and that we observe $s_1 \wedge s_2$. As can perhaps be expected, I show that the plan with the smallest cardinality is not necessarily the most probable plan. We have that,

$$P(\neg d_1, d_2, \neg d_3 \mid s_1, s_2) = \frac{P(\neg d_1)P(d_2)P(\neg d_3)}{P(s_1, s_2)}$$

and

$$P(d_1, \neg d_2, d_3 \mid s_1, s_2) = \frac{P(d_1)P(\neg d_2)P(d_3)}{P(s_1, s_2)}$$

and thus,

$$P(d_1, \neg d_2, d_3 \mid s_1, s_2) > P(\neg d_1, d_2, \neg d_3 \mid s_1, s_2)$$

if,

$$P(d_1)P(\neg d_2)P(d_3) > P(\neg d_1)P(d_2)P(\neg d_3).$$

This inequality holds, for example, when $d_1, d_3 > 1/2$ and $d_2 < 1/2$. Thus, there exist reasonable cases where the heuristics will favor a plan that is not the plan with the highest probability. Note how sensitive the result is to even slight changes in the prior probabilities of the plans.

Let us briefly reconsider Cases 1 and 2a—where the hypotheses are not mutually exclusive—using the new probability distribution specified by the Bayesian network where the basic plans are all independent.

Kautz (1987) states that if we also assume the basic plans are all equally likely, then for both Case 1 and Case 2a the most likely plan is also the one preferred by the heuristics. For example, in Case 2a we get that $P(d_2 \mid s_1 \wedge s_2) = P(d_2)$ and $P(d_1 \wedge d_3 \mid s_1 \wedge s_2) = P(d_1)P(d_3)$. Thus, if the basic plans are all equally likely, we have that $P(d_2) > P(d_1)P(d_3)$ and the heuristics favor the most probable plan.

Neufeld (1989), in his work on giving a probabilistic semantics to heuristics for preferring one plan over another, interprets the diagram shown in Figure 2 as a Bayesian network and examines which explanation in Case 2a is preferred.

Definition (Neufeld 1989): a **favours** b, if $P(b \mid a) > P(b)$.

Neufeld argues that intuition prefers d_2 as an explanation of the observation $s_1 \wedge s_2$ over the explanation $d_1 \wedge d_3$ and states that, indeed, $s_1 \wedge s_2$ favours d_2 but does not favour $d_1 \wedge d_3$. However, given the closed-world or complete-knowledge assumptions we made earlier, we can show that $s_1 \wedge s_2$ favours both explanations and thus neither is preferred. First, $s_1 \wedge s_2$ favours d_2 since, by Bayes' law and Assumption 1, we have that,

$$\begin{aligned} & P(d_2 \mid s_1 \wedge s_2) \\ &= \frac{P(s_1 \wedge s_2 \mid d_2)P(d_2)}{P(s_1 \wedge s_2)} \\ &= \frac{P(d_2)}{P(s_1 \wedge s_2)} \\ &> P(d_2), \end{aligned}$$

if $P(s_1 \wedge s_2) < 1$ and $P(d_2) > 0$. Second, $s_1 \wedge s_2$ favours $d_1 \wedge d_3$ since,

$$P(d_1 \wedge d_3 \mid s_1 \wedge s_2)$$

$$
\begin{aligned}
&= \frac{P(s_1 \wedge s_2 \mid d_1 \wedge d_3) P(d_1 \wedge d_3)}{P(s_1 \wedge s_2)} \\
&= \frac{P(d_1 \wedge d_3)}{P(s_1 \wedge s_2)} \\
&> P(d_1 \wedge d_3),
\end{aligned}
$$

if $P(s_1 \wedge s_2) < 1$ and $P(d_1 \wedge d_3) > 0$. Thus, under our closed-world assumptions, $s_1 \wedge s_2$ favours both explanations as long as $s_1 \wedge s_2$ is not a certainty and the plans d_2 and $d_1 \wedge d_3$ are not impossible. In other words, except by imposing very strong conditions, Neufeld's method does not aid us in selecting among the two plans.

Discussion and conclusion

The question arises, Why not just use probabilistic reasoning itself to come up with the best plan rather than try to give a probabilistic semantics to heuristics used in plan recognition? In a recent paper, Charniak and Goldman (1993) argue that a probabilistic approach is the only approach that is likely to succeed. A standard argument against a probabilistic approach is that too many numbers are needed and the numbers are not available. Charniak and Goldman counter that, by using Bayesian networks, only a modest set of numbers is needed and only reasonable independence assumptions are required. Let us in turn consider the assertions that the numbers are available and the independence assumptions are reasonable.

Charniak and Goldman do not counter the argument that the numbers that are required are not always available. That this is true is clear from their own examples, where at one point they give a particular probability as 1/|liquor-store|; that is, 1 divided by the number of liquor stores. (If this at first seems reasonable to you, suppose that you had to actually compute this value. Exactly which liquor stores would you count up?) A subjective probabilist would counter that probabilities are always available since probabilities are just degrees of belief. It seems though that what is wanted in some domains is to make qualitative statements such as "prefer the simplest explanation," and asking any more is to ask for false precision and sometimes meaningless numbers from a domain perspective. (Consider Case 2b where changing the probabilities by an epsilon will change the most probable plan.) Of course, in some domains frequency-based probabilities are available. For example, Bauer et al. (1993) show how to use probabilistic knowledge specific to a user, where the statistics are gathered from the domain and from observation of the user in the past in a UNIX mail application. Carberry (1990) uses a probabilistic approach to augment her focusing heuristic, (to distinguish the case where there is more than one possibility at a particular position in the context model) and in the example domain of course-advising the necessary statistics on how frequently students in particular programs take particular courses are readily available.

Now, let us consider the assumption that the basic plans are independent—as noted above, this assumption is made by Charniak and Goldman (1993) and Neufeld (1989) in their Bayesian network accounts of plan recognition and is also implied if we adopt the maximum entropy justification of the heuristics. Does it hold in every domain and if it does not, is it still a good approximation? Consider two plans that share a step such as d_2 and d_3 in our example plan library. The assumption of independence means that knowing that an agent is pursuing plan d_3 is irrelevant or unconnected to whether the agent is also pursuing plan d_2. Formally, it means that $P(d_2 \mid d_3) = P(d_2)$. This assumption does not hold in some domains. For example, suppose that there is a plaza near my house with a video store and a pizza store. My having a plan of watching a video tonight and as a result going to the plaza increases the probability that I'll also adopt a plan to eat pizza tonight, since I'm right there anyways. Thus the two plans are dependent. (This is an example of action overloading (Pollack 1992); see that reference for an extensive discussion and further references to the literature that show that humans frequently make use of this strategy.) In other domains, the assumption may hold. Unfortunately, the assumption of independence is often a difficult one to verify in the real world.

Finally, the conclusions I wish to draw are two-fold: (i) that the answer to whether a probabilistic approach or a non-probabilistic approach to plan preference is better depends on the specific domain of application, and (ii) in both approaches, we need to be aware of underlying assumptions and to test whether they are applicable in our particular domain.

Acknowledgements.

I thank Fahiem Bacchus for comments on an earlier version of this paper and for the suggestion that I look at entropy maximization as another justification of the heuristics that prefer to fit new observations into the current plan. I also wish to thank the referees for their thoughtful comments which helped me to improve the paper.

References

Allen, J. F. 1983. Recognizing intentions from natural language utterances. In Brady, M., and Berwick, R. C., eds., *Computational Models of Discourse*. MIT Press. 107–166.

Bauer, M.; Biundo, S.; Dengler, D.; Koehler, J.; and Paul, G. 1993. PHI—A logic-based tool for intelligent help systems. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, 460–466.

Calistri, R. J. 1990. *Classifying and Detecting Plan-Based Misconceptions for Robust Plan Recognition*. Ph.D. Dissertation, Brown University. Available as: Department of Computer Science Technical Report CS-90-11.

Carberry, S. 1985. *Pragmatic Modeling in Information System Interfaces*. Ph.D. Dissertation, University of Delaware. Available as: Department of Computer and Information Sciences Technical Report 86-07.

Carberry, S. 1988. Modeling the user's plans and goals. *Computational Linguistics* 14:23–27.

Carberry, S. 1990. Incorporating default inferences into plan recognition. In *Proceedings of the Eighth National Conference on Artificial Intelligence*, 471–478.

Carver, N. F.; Lesser, V. R.; and McCue, D. L. 1984. Focusing in plan recognition. In *Proceedings of the Fourth National Conference on Artificial Intelligence*, 42–48.

Charniak, E., and Goldman, R. P. 1993. A Bayesian model of plan recognition. *Artificial Intelligence* 64:53–79.

Charniak, E., and McDermott, D. V. 1985. *Introduction to Artificial Intelligence*. Addison Wesley.

Charniak, E. 1991. Bayesian networks without tears. *AI Magazine* 12(4):50–63.

Cohen, R.; Schmidt, K.; and van Beek, P. 1994. A framework for soliciting clarification from users during plan recognition. In *Proceedings of the Fourth International Conference on User Modeling*, 11–17.

Davis, E. 1990. *Representations of Commonsense Knowledge*. Morgan Kaufmann.

Kautz, H. A. 1987. *A Formal Theory of Plan Recognition*. Ph.D. Dissertation, University of Rochester. Available as: Department of Computer Science Technical Report 215.

Konolige, K., and Pollack, M. E. 1989. Ascribing plans to agents. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, 924–930.

Lin, D., and Goebel, R. 1991. A message passing algorithm for plan recognition. In *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence*, 280–285.

Litman, D. J., and Allen, J. F. 1987. A plan recognition model for subdialogues in conversations. *Cognitive Science* 11:163–200.

McKeown, K. R.; Wish, M.; and Matthews, K. 1985. Tailoring explanations for the user. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, 794–798.

Neufeld, E. 1989. Defaults and probabilities; extensions and coherence. In *Proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning*, 312–323.

Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann.

Pollack, M. E. 1986. *Inferring Domain Plans in Question-Answering*. Ph.D. Dissertation, University of Pennsylvania. Available as: SRI International Technical Note 403.

Pollack, M. E. 1992. The uses of plans. *Artificial Intelligence* 57:43–68.

Poole, D. L. 1985. On the comparison of theories: Preferring the most specific explanation. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, 144–147.

Poole, D. L. 1989. Explanation and prediction: An architecture for default and abductive reasoning. *Computational Intelligence* 5:97–110.

Sidner, C. L. 1985. Plan parsing for intended response recognition in discourse. *Computational Intelligence* 1:1–10.

Vilain, M. 1990. Getting serious about parsing plans: A grammatical analysis of plan recognition. In *Proceedings of the Eighth National Conference on Artificial Intelligence*, 190–197.