A Vocabulary for Indexing Plan Interactions and Repairs

Kristian J. Hammond*
Department of Computer Science
The University of Chicago
1100 East 58th Street
Chicago, IL 60637
hammond@cs.uchicago.edu

Colleen M. Seifert
Department of Psychology
University of Michigan
330 Packard Road
Ann Arbor, MI 48104
seifert@csmil.umich.edu

Abstract

Solving the multiple goals problem has been a major issue in Artificial Intelligence models of planning (Sussman, 1975; Sacerdoti, 1975; Wilensky, 1978; Wilensky, 1980; Wilensky, 1983; Carbonell, 1979); however, most models have assumed that the best plan for a set of goals to be satisfied in conjunction will arise from a simple combination of the best individual plans for each goal. However, human planners seem to possess an ability to look at a set of goals, and characterize them as a whole, instead of as a collection of individual goals (Hayes-Roth and Hayes-Roth, 1979). In this paper, we introduce the notion of indexing complex multiple-goal plans in terms of the *interactions* between the goals that they satisfy. We present the vocabulary requirements for representing the causality behind goal interactions, the general planning strategies used to resolve these interactions, and the specific plans based on these more general resolution strategies that are instantiated in the actual planning problem.

Indexing Plans in Memory

Solving the multiple goals problem has been a major issue in Artificial Intelligence models of planning (Sussman, 1975; Sacerdoti, 1975; Wilensky, 1978; Wilensky, 1980; Wilensky, 1983; Carbonell, 1979); however, most models have assumed that the best plan for a set of goals to be satisfied in conjunction will arise from a simple combination of the best individual plans for each goal. A problem with the "each goal first" planning theories is that they provide no vocabulary capable of characterizing goal and plan interactions in a form that allows access to past cases in memory based on these commonalities. Such goal interactions serve as critical constraints on successful plans, so that taking advantage of these constraints while selecting among and developing plans will not only produce "smarter" plans, but has the advantage of bootstrapping from plans previously developed for similar plan interaction situations.

Three basic requirements of any representational vocabulary used to describe, organize, and index plans are: first, that it characterize abstract patterns of goal interactions that capture relevant similarities between situations; second, that it provide access to general strategies that pertain to resolving the overall goal/plan situation; and finally, that it identifies specific plans that cover the current situation. In the next sections, we present a representational vocabulary that characterizes the causal knowledge behind goals, plans, and their interactions. Human experimental evidence is then presented, along with suggestions about how this proposed paradigm can be extended to encompass a majority of planning situations.

Vocabulary for Goal Interactions

Planners currently use a vocabulary of goals, associated plans, sub- plans, preconditions, and effects (Schank and Abelson, 1977), as well as basic interactions such as *conflict* and *concord* (Wilensky, 1978). The problem of how to describe the similarity between goal situations has been discussed by (Schank, 1982), who introduced abstract memory structures (Thematic Organization Packages or TOPs) to connect episodes in memory on the basis of similarities in the pattern of goals and plans they contain. In planning, such abstract patterns of goal and plan interactions can serve to identify a class of problems where a particular set of resolution strategies are appropriate.

Consider the example (Wilensky, 1978) of "wanting the newspaper from outside on the sidewalk

*This work was supported in part by AFOSR grant number AFOSR-91-0112, DARPA contract number F30602-91-C-0028 monitored by Rome Laboratories, DARPA contract number N00014-91-J-4092 monitored by the Office of Naval Research, Office of Naval Research grant number N00014-91-J-1185

while it is raining," where the planner is trying to achieve a particular goal (getting the newspaper); the chosen plan (carry the paper in) has a particular precondition (be outside); and an existing state (it is raining), in combination with the precondition, results in the violation of an existing preservation goal (stay dry) (Schank and Abelson, 1977). To plan in this situation, the goal conflict must be described in terms of an abstract characterization of the problem that captures the causal chain leading to the violation. This situation can be characterized as *precondition* plus *state* causes *violation* of a preservation goal, or *Plan+State→Violation* (see Figure 1).

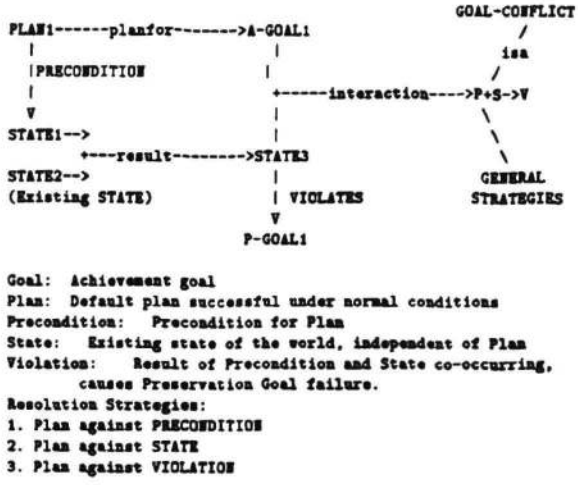

Figure 1: Representation of the goal interaction situation Plan+State→Violation.

Many other plan/goal conflicts can be characterized in a similar way (see (Hammond *et al.*, 1991; Hammond, 1990)). In this paper, we will concentrate on only one of these, *Plan+State→Violation*, and its associated resolution strategies.

Vocabulary for Resolution Strategies

There are three basic strategies that are designed for use in any *Plan+State→Violation* planning situation. These resolution strategies are plans for altering the causal situation such that the original goal can be achieved while avoiding the preservation goal violation. The three strategies for *Plan+State→Violation* are:

- Plan against precondition: Find a plan for the primary goal that does not require the problematic precondition (find a plan that does not require going outside).
- Plan against state: Alter the existing state such that even when the precondition is met the preservation goal will not be violated (do something to make it stop raining).
- Plan against violation: Add an auxiliary plan that prevents the violation of the preservation goal even in presence of the threatening state and the existing condition (get an umbrella).

One test of these problem resolution strategies is whether or not they apply to other instances of *Plan+State→Violation*. Suppose a planner wants to get a hot pot off of the stove. Like the newspaper and the rain example, the situation fits into *Plan+State→Violation*, and the associated resolution strategies are appropriate here also: the first strategy suggests trying a plan that does not require touching the pot, such as using a stick to push the pot off of the burner; the second suggests waiting for the pot to cool down before moving it; the third suggests that some sort of protection be used, such as a pot holder. While the content of the particular goals is different, the interactions that occur between the goals in both of these situations are similar.

Accessing general strategies through this vocabulary provides the planner with alteration techniques and information as to what parts of the initial causal configuration are appropriate targets of change. In summary, by including a causal analysis of the goal interactions in a situation as part of its representation, it is possible to access in memory the general strategies applicable to the problem, leading to specific plans for the current situation.

Vocabulary for Specific Planning Strategies

Causal relations can be used to organize resolution strategies in terms of the situations for which they are relevant. The resolution strategies indicate *where* a particular causal chain can be effectively altered for a particular planning situation; next we need a more specific vocabulary for characterizing *how* situations can be altered. We will now look at when and how to apply the specific strategies associated with *Plan+State→Violation*.

Specializations of "Plan against precondition." One specialized strategy is *use alternate agent*, appropriate in *Plan+State→Violation* cases such as when a student wants to go into the office and pick up his mail while avoiding his advisor. This strategy is not always appropriate; for example, consider moving the hot pot. To index this strategy so that it will be applied only in appropriate problems, we need a representation of features that identify when this and other particular strategies are relevant. In this case, the distinctive feature lies in the nature of the preservation goal goal being threatened: if it is specific to the planner and not to other agents, then this is a good solution to a *Plan+State→Violation* conflict. Thus, the

important features for *use alternate agent* are the commonality of the threatened preservation goal, the nature of the preservation goal, and any special skills or abilities involved in the normal plan to satisfy the initial achievement goal (Schank and Abelson, 1977).

A different specialization of the general strategy to plan against the precondition is to use an alternate plan that does not have the same precondition, or *use alternate achievement goal plan*. For example, to get the hot pot off of the stove, one could pick up the pot without touching it by inserting a stick through its hollow handle. A predictive feature (Johnson and Seifert, 1991) for this strategy is the existence of any alternate plans for the initial goal.

A final specialization of this general strategy of planning against the precondition is to run the initial plan very quickly, the *run fast* strategy. This strategy is effective in those cases where the preservation goal being threatened has degrees of violation linked to a parameter (time or speed) under the agent's control, and when the preservation goal violation is only intermittent (as in the possibility of running into the advisor in the mail room).

Specializations of "Plan against state." As with "plan against precondition," there are three specializations of "plan against state:" *wait out, jump between, and counter plan. Waiting out* the state applied to instances of *Plan+State→Violation* in which the existing state is temporary, such as a hot pot. Use of this strategy depends on the nature of the achievement goal as well as on the nature of the state; if the achievement goal is particularly insistent, then this strategy is inapplicable.

A variation on the *wait it out* strategy is the *run between* strategy, where one runs the initial plan in between fluctuations in the state. This strategy would be indexed under the *Plan+State→Violation* structure, linked to a test concerning the permanence and possible intermittence of the side state being planned against.

Finally, one can *counter plan* against the conflicting state; that is, generate a separate plan to disable the state. For this to work, the planner obviously must have some control over the state. Once again, the indexing vocabulary for all three strategies consists of features related to the practical considerations that the planner has to take into account in order to alter the initial plan: duration of states and existence of specific plans.

Specializations of "Plan against violation." Like the other two general strategies, "plan against violation" has some specialized versions that can be applied in different circumstances. If the preservation goal is relatively minor or short-lived in relation to the achievement goal, then it might make sense to *tolerate* violations of it. To decide if *toleration* is possible, the planner needs to know the relative importance of its goals, and the likely durations of violations.

Like *tolerate*, *run and recover* requires not only that the preservation goal is relatively unimportant, but also that there exists a plan associated with the violation that can be used to recover from it. One can, for example, dry off after fetching the newspaper out of the rain. Like the *tolerate* strategy, this *run and recover* strategy depends on the relative importance of the two goals being planned for, and requires a test for the existence of recovery plans for the violated goal.

A third specialization of "plan against violation" is *counter plan the preservation goal*: generate a concurrent plan in support of the preservation goal, such as using an umbrella. The plan is stored in terms of the causal situation in which it will become relevant, rather than in terms of a specific goal violation.

The specializations of the general strategies apply in some instances of *Plan+State→Violation*, but not others, and therefore must be indexed by their appropriateness conditions. That is, "plan against violation" has three specializations *tolerate* applies when the preservation goal being protected is trivial compared to the achievement goal being satisfied; *run and recover* applies when there is a straightforward recovery plan associated with the violation; and *counter plan the preservation goal* applies when the agent has an existing plan associated with the causal rule leading to the violation, and also when an additional state is required for the violation to occur (such as physical contact).

To summarize this vocabulary, when a situation where a precondition and an existing state cause a preservation goal to be violated, we can respond with one of three possible resolution strategies: plan against precondition, plan against state, and plan against violation. Each of these strategies has several specializations whose appropriateness depends on the pragmatic planning constraints in the situation.

Indices for Retrieving Planning Cases

A plan is proposed, a conflict detected and characterized, a specific strategy chosen. The next step is to search memory for a past instance of that strategy that applies to the current situation; however, we must define the set of indices by which those instances can be recalled at the appropriate time. Every instance of a particular strategy is indexed by

the features of that episode that the strategy used or altered to construct the instance, and by implementational causes and effects that are learned through experience. In general, then, the features used to index planning cases are those which have some causal relevance to the way in which that strategy is implemented.

For each of the nine planning strategies proposed above, the following features are likely to lead to useful past instances:

- Indices for *use alternate agent* include: The achievement goal, other available agents, the plan itself, the threatened preservation goal, and the state threatening the preservation goal.
- Indices for *use alternate plan* include: The achievement goal, the specific plan, possible alternate plans, and the precondition to be avoided.
- Indices for *run fast* include: The plan, the state and precondition implicated in the violation of the preservation goal, and the preservation goal being violated.
- Indices for *wait it out* include: The achievement goal and the state
- Indices for *Jump between* include: The achievement goal, the proposed plan, and the intermittent state.
- Indices for *counter-plan state* include: The state, the available plans for that state, and the harmful precondition of the initial plan.
- Indices for *tolerate the violation* include: The preservation goal and the violating state.
- Indices for *run and recover* include: The violating state and possible plans to recover from the violation.
- Indices for *counter-plan violation* include: The rule connecting the precondition and the undesired state, the assumed conditions in that rule, the undesired state, the precondition for the initial plan, and the existing state in the world.

We have presented a set of specific strategies that can be applied in different situations described by *Plan+State→Violation*, each of which has a question or feature that tests for its applicability. The structure *Plan+State→Violation*, then, includes more than just how, in general, to alter the resolution strategies: It also has information about how to *apply* specific planning strategies, and in what circumstances the individual strategies are relevant. The features useful for determining the applicability of these strategies are critical to this new vocabulary for describing planning situations.

Evidence for Vocabulary Use

To determine whether humans utilize causal planning factors in selecting among planning strategies, we conducted an experiment in common-sense planning (see (Hammond *et al.*, 1991), for a full description). The planning problems used were six exemplars of the *Plan+State→Violation* structure, all placed in different contexts such as celebrating a sick friend's birthday, jogging after dark, and picking up an exam while avoiding one's professor. Subjects were asked to provide commonsense answers to the planning problems in terms that they would really chose to do in those situations. By examining the types of plans they propose, and how well those plans match the predictions from the vocabulary model, it can be ascertained whether subjects are utilizing these same features in determining plan selection. The responses were then coded using the planning strategies predicted by the vocabulary model. Any responses not fitting one of the categories was coded in a general "other" category. These were:

1. Plan against Precondition: Find a new plan for achieving the primary Goal that does not require the Precondition which threatens the preservation Goal.
   (a) Get alternate agent (if not problematic for other actor)
   (b) Run Plan fast (if limited exposure is acceptable)
   (c) Use alternate plan without Precondition (if available)
2. Plan against Existing State: Alter the Existing State so that even when the Precondition is met, the Preservation Goal will not be violated.
   (a) Waitout Existing-State (if temporary)
   (b) Jump in between phases of State (if intermittent)
   (c) Counterplan against Existing-State (if possible)
3. Plan against Violation (Threatens Preservation-Goal): Add an auxiliary plan that prevents the violation of the preservation Goal even in the presence of Precondition and State.
   (a) Ignore – put up with preservation Goal Violation (if short duration)
   (b) Plan to recover from preservation Goal Violation (if can repair)
   (c) Counterplan Preservation Goal Violation (interrupt connection between Precondition and State)

Overall, the extent to which the responses given fit into the proposed categories support the use of the causal features in commonsense planning, as

opposed to other features or plans subjects may potentially generate for the problems.

93% of responses were instances of the planning strategies proposed in the vocabulary theory, while 7% were "other" types of responses. The "other" responses included items such as "see if the dorm has anyone else around and borrow their newspaper" for the newspaper in the rain problem, or "stay home" for the running after dark example. The majority of "other" responses involved abandoning the goal implicated in the goal interaction. This type of response is not predicted by the vocabulary model, which assumes the goal must be satisfied in some way. Overall, it appears the planning strategies for the *Plan+State→Violation* structure were sufficient to account for the plans generated by subjects, with the exception of solutions involving abandoning goal satisfaction.

Among the three resolution strategies, subjects' responses more frequently involved planning against the violation (43%), compared to planning against the precondition (33%) or planning against the existing state (17%). While the model makes no predictions about the use of the three categories beyond which features apply in specific instances, it seems subjects preferred plans that dealt directly with the problematic interaction of precondition and state, rather than attempting to change either separately. In particular, plans to change or work around the existing state of the world were given infrequently compared to other possibilities. This may reflect task demand, in the sense that subjects tried to work within the problem constraints presented, and the states tended to be examples of conditions in the world that are unchangable (such as rain and darkness).

In addition, comparisons of strategies by example indicates high variability in strategy application based on the specific planning constraints in each of the examples. The results show that, while each strategy was used in at least 3% of responses, an uneven pattern of strategy use across examples was evident. Of the specific strategy instantiations, the most frequent was the strategy of "counterplanning against the preservation goal violation," with 27% of responses. Another example of selective use of specific strategy is "selecting an alternate plan without the problematic precondition." This strategy, applied only when such a plan is available, was frequently used in the "driving to Detroit" example, where substituting other means of transportation avoids the faulty brakes in the planner's car. For three of the problems, no responses included substitute plans that avoided the precondition.

In addition, each planning problem showed distinct differences in application of resolution strategies. For example, the most frequently generated plan for the newspaper example was "counterplan against the violation," while for the "driving to Detroit" example, the most frequent plan was "use alternate plan." The reason for these differences rests in the pragmatic features used to determine when a plan is appropriate for application. For example, the use of an alternate plan depends on the existence of such a plan, most obvious in the "driving to Detroit" example where other means of transportation are readily available. Subjects did not perceive many alternate plan for getting the newspaper not involving going outside. Thus, there were many differences in the patterns of plans generated for each example. In general, these patterns fit the predicted categories, such that no plans were generated when the vocabulary tests suggested that the planning strategy was not appropriate. However, there were instances where subjects did not generate plans that could have been expected based on the tests in the planning vocabulary. For example, the strategy of "waiting until the rain lets up" was predicted for "jumping in between" intermittent phases of rain in the newspaper example, but not listed by subjects. However, subjects may have felt information about the state of the rain was lacking, and so avoided using plans based on assumptions about the state not given explicitly in the problem description.

In conclusion, it appears that the proposed planning vocabulary accounts for the set of responses given by subjects to these simple planning problems. Further, there was good evidence that subjects were sensitive to the applicability features associated with each strategy, such that they applied some strategies only in appropriate examples. The vocabulary did not include any prediction of the demonstrated preference for plans against the violation itself, compared to plans against either the precondition or the state separately. The representational scheme also did not account for unsuccessful plan resolution, whereas most of the "other" responses involved subjects' attempts to abandon the goal. Overall, however, the plans subjects generated corresponded extremely well to the causal possibilities laid out in the vocabulary, and few novel intrusions occurred. Further, the application of strategies differentially in the specific problems supports the notion that subjects are sensitive to the features predicting when certain strategies are applicable.

Generality of Vocabulary

The *Plan+State→Violation* vocabulary includes many features that are important for planning in general. It is clearly important for a planner to to know the difference between those states that it can plan against and those that it cannot; what the

preconditions and effects of its plans are; if it has any other plans for the same goal that has different preconditions and effects; and if the goals it has are only held by it, or also held by others with whom it can share tasks.

The indexing within *Plan+State→Violation* uses the same features a planner needs to detect and monitor in order to plan, and are neither arbitrary nor important only to this structure. Thus, in *Plan+State→Violation*, nine specific strategies are stored under the TOP, indexed by the features of the goals, actions, and states in the structure that determine the applicability of the strategies themselves. The components of structures representing goal interactions should include those of *Plan+State→Violation* (preconditions, existing states, violations) along with many others to describe the prototypical ways in which goals can affect one another (side effect, disable, enable, etc.) Similarly. the set of resolution strategies outline for *Plan+State→Violation* must be extended to capture different modifications to other goal interaction structures. Finally, the specific plan strategies for each situation are greatly affected by the context, and will vary based on the specific features of the planning problems.

With the aid of a complete vocabulary of plan/goal interactions, the planner can, after identifying its situation as an instance of a particular causal structure, apply a few simple tests to select from a set of easily implemented strategies. A refined understanding of the causal pattern underlying the prototypical solutions also allows the generation of alternate solutions when needed; for example, when an umbrella is not available, other materials can serve the same functional purpose.

Conclusion

This paper has presented an outline of a representational scheme for organizing and accessing plans and past episodes relevant to current planning problems. Our argument is that it is the abstract relationship among goals and plans that best constrains what planning choices one might make in a given situation. Therefore, retrieving past plans based on the abstract interaction will assure that the retrieved information will be most useful to the planner.

For any particular planning problem, this representation allows easy access to general resolution strategies and specific past plans related to the goal interaction situation. This allows the planner to search for past plans relevant to its overall situation, rather than building complex plans out of single plans for each of the goals in its current planning situation.

The vocabulary required to support this organization connects three levels of abstraction in planning. In order to identify the particular TOP relevant to any given situation, the planner must be able to characterize its current goal/plan problem in terms of the causal relations between the goals, actions and states included in that episode. This characterization then allows the planner to identify the TOP which packages the general strategies applicable to its current problem. To select among the general strategies, the planner must answer pragmatic questions about its current goal/plan configuration.

Among all possible features in a planning situation, only a limited set of these features – those that are relevant to the way in which the current causal structure can be changed – are used as indices within the TOP.

References

J. Carbonell. *Subjective Understanding: Computer Models of Belief Systems.* PhD thesis, Yale University, 1979.

K. J. Hammond, C. Seifert, and K. Gray. Intelligent encoding of cases in analogical transfer: A hard match is good to find. *The Journal of the Learning Sciences*, January 1991.

K. J. Hammond. Explaining and repairing plans that fail. *Artificial Intelligence Journal*, 1990.

B. Hayes-Roth and F. Hayes-Roth. A cognitive model of planning. *Cognitive Science*, 3(4), 1979.

E. Sacerdoti. A structure for plans and behavior. Technical Report 109, SRI Artificial Intelligence Center, 1975.

R. C. Schank and R. Abelson. *Scripts, Plans, Goals and Understanding.* Lawrence Erlbaum Associates, Hillsdale, New Jersey, 1977.

R. Schank. *Dynamic Memory: A Theory of Reminding and Learning in Computers and People.* Cambridge University Press, 1982.

G. Sussman. *A computer model of skill acquisition*, volume 1 of *Artificial Intelligence Series.* American Elsevier, New York, 1975.

R. Wilensky. *Understanding Goal-Based Stories.* PhD thesis, Yale University, 1978. Technical Report 140.

R. Wilensky. Meta-planning. Technical Report M80 33, UCB College of Engineering, August 1980.

R. Wilensky. *Planning and Understanding: A Computational Approach to Human Reasoning.* Addison-Wesley Publishing Company, Reading, Massachusetts, Reading, Mass, 1983.