The Operational Level of a Commonsense Planner

Richard Alterman
Department of Computer Science
Brandeis University

ABSTRACT

This paper characterizes the operational level of a commonsense planner. In particular it will compare two approaches on the problem of level of operation. A 'plan instantiater' planner has access to a sort of summarization of the activities involved in each of the subkinds of a given category of plans, and by a process of instantiation it selectively adds details to match the current planning situation. Such planners abandon the details of the lower plan and dynamically recreate them under the exigencies of a given situation. A 'reference point' planner selects a subordinate plan, from a given category of plans, to represent the category as a whole. The reference point planner assumes a level of operation which can directly access a greater number of functional details than its plan instantiation counterpart, but perhaps at the loss of flexibility.

The thrust of this paper is that reference point planners are the appropriate model of planning for the commonsense domain.

1. Introduction

The term commonsense planning refers to the mundane day-to-day activities of human planners [21]. This includes the planning of such activities as fetching a newspaper, making a deposit at the bank, taking a child to daycare, fixing breakfast, riding a bus to work, going to a movie, making an airline reservation, shopping at a supermarket, eating at a restaurant - in short - planning for the quotidian concerns of human existence as it varies under the flux of daily circumstances. In commonsense planning it is less the case that the activities of the planner vary a great deal, and more the case that the circumstances underwhich the plans are applied vary.

Adaptive Planning [1,2] is an approach to planning in the commonsense domain. An adaptive planner takes advantage of the habitual nature of many of the planning situations for which it plans by basing its activities on pre-existing plans. A critical issue, and the subject of previous papers, is the question of flexibility: How does an adaptive planner refit an old plan in order to meet the demands of some new planning situation?

One of the important themes of adaptive planning is the way it characterizes the level of operation. It takes the following position: specific plans are foregrounded (in order to take advantage of the mass of details associated with the more specific plan), and general plans are in the background (serving more to organize old plans into categories). This paper is intented to act as a clarification of this position. In particular it will compare two approaches on the problem of level of operation. Approach one is that commonsense planners work from plans that are basically underspecified and then add details as the situation demands. Approach two is that the planner works from plans that, if anything, are overspecified and then removes details and refits the plan only as indicated by the situation. A critical point in this discussion will be the habitual nature of the commonsense planning domain. The paper is fairly tendentious (in favor of the second approach).

Consider the category of plans concerning TELEPHONING (see figure 1).

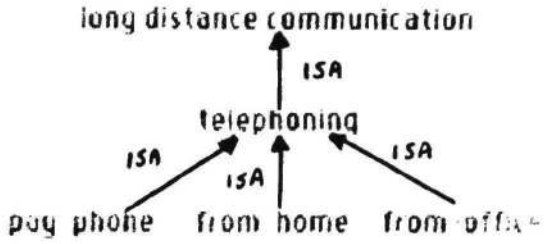

Figure 1: Two Approaches.

I will refer to a planner that works at the level of TELEPHONING as a planner that works via a process of **plan instantiation**. The notion here is that the planner has access to a sort of summarization [17] of the activities involved in each of the subkinds of TELEPHONING, and by a process of instantiation it selectively adds details to match the current planning situation. Such planners abandon the details of the lower plan and dynamically recreate them under the exigencies of a given situation.

I will refer to a planner that works from one of the subordinate level plans (i.e. from-office, from-home, or pay-phone) as a **reference point** [10,14] planner. Each of the subordinate plans listed for TELEPHONING can serve as a reference point for the category depending on the planning context. The general idea here is to use a subordinate plan to represent the category as a whole. The reference point planner assumes a level of operation which can directly accesses a greater number of functional details than its plan instantiation counterpart, but perhaps at the loss of flexibility.

This paper will articulate the trade-off between a planner that works operationally as a plan instantiator and one that works operationally as a reference point planner.

2. Generalizing the Information in a Subordinate Plan

I will distinguish between two sorts of generalization: object generalization and sequence of action generalization. Object generalization is concerned with generalizing the objects, or props, associated with a given set of plans: an object that can be generalized out of the three telephone plans is the 'telephone'. A sequence of action generalization is concerned with generalizing over the steps of the subordinate plans. The focus of this discussion will be on sequence of action generalizations: I will give a detailed accounting of the kinds of action knowledge that is lost as a result of summarizing a given category of plans. This analysis should support two sorts of conclusions. First that making a useful generalization for a given set of subordinate level plans is by-and-large a difficult proposition. Second that even in the event that a useful generalization can be made, the problem of re-instantiating details on the fly for a given habitual situation is doubly complicated since the generalization has lost not only the details associated with individual steps, but also the details that correspond to the inter-relationships amongst those steps.

2.1 Delete Steps

Below are shown the sequence of actions associated with the two kinds of TELEPHONING:

from-home - 1) lift receiver 2) hear dial tone 3) remember number 4) dial 5) ringing sound 6) ask-for 7) talk
pay-telephone - 1)get change from pocket 2) lift receiver 3) hear dial tone 4) insert coins 5) dial number 6) ringing sound 7) ask-for 8) talk

It would be possible to form a generalization over these two kinds of telephone calling by removing two steps from the pay-telephone plan: getting change from one's pocket and inserting coins into the telephone box. There is an implicit loss of default information when individual steps are deleted. The subordinate pay telephone plan encodes the fact that the planner's preferred plan for getting change is to reach into his/her pocket. For another plan, say doing the laundry at the laundromat, the preferred plan for getting change is to use a change machine.

2.2 Generalize Individual Actions

Consider the following two plans:

theatre-plan - 1) buy theatre ticket 2) enter playhouse 3) view play 4) leave
movie-plan - 1) buy movie ticket 2) enter movie 3) view movie 4) leave

In this case it would be possible to form a generalization by generalizing three of the individual steps, i.e.,

show-plan - 1) buy ticket 2) enter show 3) view show 4) leave show

There is a cost involved in generalizing an individual step. Consider the first steps of the theatre and movie plans. They both involve buying tickets, but one usually requires a reservation and for the other a reservation is not possible. In the theatre plan situation the planner calls and makes a reservation, thus saving a trip if the show is sold out. Similar sorts of loss of information occur for step two (in one case refreshments are bought before the show) and step three (in one case there is an intermission). Where a plan instantiator must re-discover these details in a piecemeal fashion, a reference point planner, by selecting a subordinate level plan, accesses all of the details, of one or the other plan, all at once.

2.3 Relax Sequence Constraints

There is another kind of information that can be lost when generalizing over subcategories of a plan and that is sequencing information. Consider three subcategories of a restaurant plan:

sitdown restaurant - 1) seating 2) ordering 3) serving 4) eating 5) paying
fastfood plan - 1) order 2) pay 3) serve 4) sit 5) eat
cafeteria plan - 1) select food 2) pay 3) sit 4) eat

These three plans include many of the same actions, but they are performed in a differing orders. In this case, it would be possible to factor out many of the actions, with some generalizations, but sequencing information would be lost.

Relevant to this issue is the problem of sequencing. There are three sorts of positions taken on the sequencing problem in the AI literature. One approach (linear assumption [19]) is that the planner commit itself to any ordering of actions, and if difficulties arise, due to subgoal interactions, make the appropriate changes. The second approach (least commitment [15]) is to make a commitment to an ordering only when necessary. The third approach (opportunism [6]) is to perform a given step when convenient. Working from reference points offers a fourth approach: Select an old ordering. The set of reference points, in effect, offers a short list of possible orderings. Encoded in each reference point are sequencing considerations, many of which are dictated for causal sorts of reasons. A good heuristic for domains of habitual activity is to select, or reselect, orderings from the short list.

3. Does a good reference point always exist?

How do we know that there is always a reasonable reference point? Suppose the category of plan concerned with telephoning consists of the following two subordinate plans:

> **from-home** - 1) lift receiver 2) hear dial tone 3) remember number 4) dial 5) ringing sound 6) ask-for 7) talk
> **pay-telephone** - 1)get change from pocket 2) lift receiver 3) hear dial tone 4) insert coins 5) dial number 6) ringing sound 7) ask-for 8) talk

Furthermore suppose the new situation is that the planner intends to make a telephone call from work and must dial a nine in order to get an outside line. Claim: if there exists a generalization over these two plans that is applicable and reasonable for the planning situation above, then one of the two reference points is reasonable for the same planning situation. Or, more generally, if for a given planning situation there is a reasonable pre-existing general plan, then there always exists a reasonable reference point plan for the same situation. Below is shown an argument as to why this must be the case.

[1] If a generalization gets to be too abstract then it stops being useful. Consider the following example: Suppose a graduate student comes to you and asks you for a plan to get a Phd. One answer is: pass your generals, pass your orals, write a dissertation, defend your dissertation. Such a plan is a description of what the student must do, but it is much too abstract to be of any use.

[2] If a plan P is instantiated from a generalization G, G is a generalization of P.

[3] Given a set of plans the amount of abstraction that is required over that set in order to form a generalization works as a function of the difference between any pair of plans in that set. The greater the difference between any two plans in the set, the more abstract the general plan must be. Given any number of plans $P_1 ... P_n$ there exists a reasonable generalization over these plans iff the differences between any two plans, $\Delta(P_i, P_j)$, is reasonable as defined by some constant C, i.e.

$$\Delta(P_i, P_j) < C \quad for\ 1 \leq i, j \leq n$$

[4] Any plan P which is derived from a generalization G is part of the set of plans which G generalizes over (2), and consequently, by our notion of a more general plan, must not be too different from the other plans in the class (3), i.e.

$$\Delta(P_i, P) < C \quad for\ 1 \leq i \leq n.$$

The above shows that if there is a reasonable generalization which applies to a given situation then there is a reasonable subordinate plan which also applies. That does not mean that there is always a reasonable pre-existing general plan. The previous section demonstrates some of the difficulties in forming a reasonable general plan. In general, it is not hard to find cases where there exists a category of plans where it is difficult to find a useful summary, but is nevertheless a useful category of plans. Consider the following two plans for telephoning.

from-home 1) lift receiver 2) hear dial tone 3) dial 5) ringing sound 6) ask for 7) talk
from-work 1) lift receiver 2) buzz secretary 3) request call 4) hang-up and wait 5) hear buzz and pick-up receiver 6) greet and talk

It is hard to see what a useful generalization would be over these two plans, yet it is clear that they represent two kinds of telephoning. Moreover, there are any number of planning situations where one of these two subordinate level plans would be an appropriate starting point.

Perhaps it would be possible to get around this kind of problem by splitting the category or allowing for multiple generalizations for a given category of plans. These kinds of techniques would work - the generalizations are making finer and finer discriminations - but dividing up the category is just another way of converging on reference points.

There is an assumption here that the planner is working from prestored plans. Perhaps if the more general plans are dynamically created it would be possible to work from more general plans. A whole generation of AI planners, beginning with [15], has been constructed along these lines. But there are difficulties for this sort of approach in the commonsense domain. Largely these difficulties arise from the fact that the planner fails to take advantage of the habitual nature of most of his/her activities. The constraint of habituality decidedly favors using pre-existing plans.

Notice that there is no claim here that there always exists a useful pre-existing plan. There will be planning situations that require the planner to plan for a non-habitual activity. The point is that for planning and acting with regards to habitual activies there is nothing to lose - and a lot to gain - in selecting and working from a reference point.

Finally, the claim is about the existence of a reasonable reference point plan. There is a question as to how hard, or easy, it is to find a reference point plan. The associative techniques suggested by [18,7] are good candidates for the task of finding a reference point for a given planning situation. Since the activities we are concerned with are habitual, they should be relatively easy to find.

4. Flexibility and Levels

How can a commonsense planner maintain the advantages of both flexibility and habituality? One candidate answer is to have the commonsense planner work from specific plans and when things go awry, it can retreat to weak methods [3,4]. Under certain circumstances this could be the appropriate choice, but there are many instances where the planner can succeed without a full retreat to weak methods. In this section I will describe a reference point planner PLEXUS that can perform flexibly without retreating to weak methods. For further details and traces see [1,2].

Assume that the reference point planner has access to the content and organization of background knowledge associated with the old plan. In [1,2] the importance of background knowledge is described, as well as some details about four kinds of background knowledge. Here, it will suffice to draw attention to one of the kinds of background knowledge, distributed categorization knowledge. Distributed categorization knowledge refers to the categorization hierarchies associated with the various steps and subplans of the reference point plans (as opposed the categorization hierarchy associated with the overall plan).

Assume the commonsense planner selects a reference point from the category of plans that it intends to accomplish. In a given situation, if a problem arises, the planner's refitting behavior is directed towards finding an alternate reference point. Alternate reference points can take one of three forms:

[1] *Select a different reference point from the overall category of plans.* Suppose the planner intends to make a call from work, and its usual plan for the office situation is with-secretary (see figure 2). Furthermore, suppose the secretary was out of the office. In this case planner is directed towards choosing an alternate reference point for the overall category of plan, i.e. from-home.

[2] *Select a different reference point for the individual step.* Suppose the planner decides to use a pay-phone and discovers, as it attempts the first step of the plan, that it has insufficient pocket-change (see figure 3). Adapting to this situation does not require that the overall plan be changed - instead an alternate version of the individual step must be used. Within the distributed categorization hierarchy associated with getting-change there exists several alternate reference points: getting change from a merchant, asking somebody else for change, or using a change machine. Depending on the current planning circumstance one or another of the reference points is appropriate.

[3] *Select a different reference point for the implicit subplan.* A given plan can be construed as a collection of interleaved subplans. Under certain circumstances the planner must find an alternate reference point for one of the subplans. Suppose the planner intends to make a telephone call from a pay-phone and discovers that it has no money at all (see figure 4). In this case it is the implicit pay-direct subplan for which an alternate reference point needs to be found, i.e. call-collect

The commonality for each of these methods is that they all involve **abstraction** and **specialization** within a categorization hierarchy. By a process of abstraction the planner removes details that are not appropriate to the planner's current circumstances. By a process of specialization the planner adds details, *en masse*, which are appropriate for the

current circumstances. See [1,2] for heuristics that apply to each of these processes and a description of four kinds of situation difference that can occur between an old plan and a new situation.

Each of the above methods of adaptation require that the planner return to the reference point level of operation. In a sense the higher levels of the categorization hierarchies act in an organizational capacity, while the more special plans capture the habitual qualities of the commonsense domain.

5. Discussion and Summary

5.1 Discussion

Until recently, Wilensky's work on commonsense planning virtually stood alone in the AI literature as a computational model of a commonsense planner [21]. As oppposed to weak method reasoners [12] that apply to knowledge poor domains, Wilensky contended that commonsense planning was knowledge intensive. Although the emphasis of his work was on developing goal and plan meta-structures that were shared by both planning and understanding activities, Wilensky did propose an architecture for commonsense planning based on a cycle of plan proposal, plan projection (or simulation), and goal detection. Wilensky assumed that his planner was working from the most specific plan that was appropriate to a given situation, but never really developed or defended this idea.

The notion of cognitive reference points is attributable to Rosch [14]. This paper has been most directly influenced by the recent work of Lakoff on reference points (Lakoff refers to reference point reasoning as metonymic reasoning) [10]. Both Rosch and Lakoff were interest in using reference points to account for prototype effects. This work complements these previous studies by looking at domains other than natural kinds and natural language (i.e. plans). Here the role of the reference point is important because it encapsulates many of the important details that are not available at higher levels.

This work is also related to the recent work on case-based reasoners [8,9,13,5]. For the purposes of this paper I would like to differentiate between two types of knowledge that case-base reasoners can access: semantic knowledge and episodic knowledge. Roughly, an episode is a particular experience in the biography of an individual, and semantic knowledge refers to the conceptual thesaurus of the individual [20]. Semantic knowledge can be used to index and interprete the episodes of an individual [16]. Episodes can be interpreted by more than one piece of semantic knowledge. When a case-base reasoner works from an episode, it must first extract the relevant plan structure and then adapt it. When a case-base reasoner works from semantic structures, it need only adapt the structure. By-and-large the habitual nature of the commonsense domain suggests that the planner is working from semantic structures - in all likelihood the correct interpetation of an episode has already been made.

5.2 Summary

A plan instantiator planner works at the level of the general plan and recreates the details of a more specific plan as the situation demands. A reference point planner selects one of the subordinate plans as a basis for planning activity, thus it can access a greater

number of functional details. This paper presents several analyses which suggest that commonsense planners are reference point planners.

One analysis demonstrated the difficulties associated with forming generalizations over sets of plans. In order to create a general plan from a subordinate plan, three kinds of changes can be necessary: the deletion of steps, the generalization of the individual actions of the plan, and the relaxation of sequence constraints. Each kind of change represents a loss of information and has an associated cost.

Whenever there exists a reasonable pre-formed general plan for a given task, there exists a reasonable reference point plan. In many cases the subordinate plans differ sufficiently to make the formation of a useful generalization difficult, if not impossible - nevertheless the subordinate level plans of such categories can prove to be quite useful. Planners that dynamically create general plans, before instantiating them, are, on the whole, inappropriate for commonsense planning because they fail to take advantage of the habitual nature of the commonsense domain.

A reference point planner can achieve flexibility by taking advantage of the distributed categorization hierarchies that exist in the background knowledge. Achieving flexibility is largely a function of finding an alternate reference point. PLEXUS is an example of a commonsense planner that adapts a plan, without resorting to weak methods, via the selection of alternate reference points. In a given situation PLEXUS can adapt a plan in one of three ways: it can select a different reference point from the overall category of the plan; it can select a different reference point for an individual step; it can select a different reference point for the implicit subplan. Each of these three adaptive methods can be characterized as a process of abstraction and specialization.

In some ways this argument is analogous to the following choice of wagers:

Heads you win. Tails no bet.

Heads you win. Tails you lose.

For habitual activites reference point planning is equivalent to the first wager and plan instantiation is equivalent to the second wager. Under the best of circumstances the reference point plan is tailor-made for a given situation, and under the worst of conditions it is no worse than working from a general plan via a process of plan instantiation. Thus in domains of habitual activity - that is domains where habitual activities are the rule and not the exception - gearing the planner, and its representation of planning knowledge, towards reference points can be a saving heuristic.

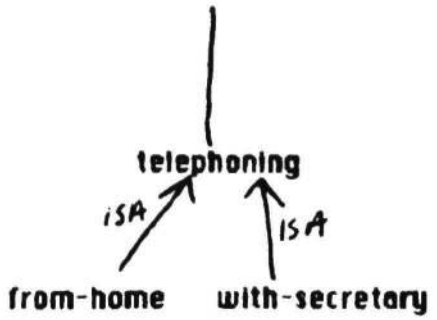

Figure 2: Overall Plan

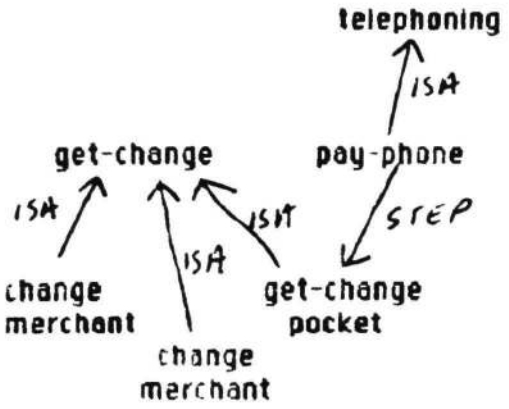

Figure 3: Individual Step.

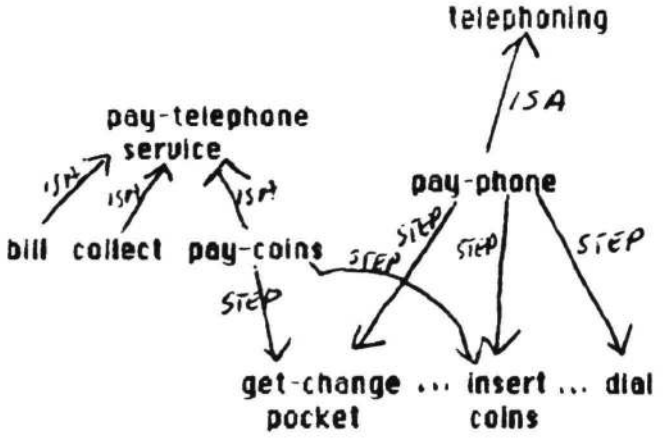

Figure 4: Implicit Subplan

REFERENCES

[1] Alterman, R. "An Adaptive Planner" in AAAI-86.

[2] Alterman, R. "Issues in Adaptive Planning." Technical Report No. UCB/CSD 87/304, Computer Science Division, UC Berkeley, 1986.

[3] Carbonell, J. "A computational model of analogical problem solving," in IJCAI 7, 1981.

[4] Carbonell, J. "Derivational analogy and its role in problem solving," AAAI-83.

[5] Hammond, K. "Chef: A model of case-based planning" in AAAI-86.

[6] Hayes-Roth, B and Hayes-Roth, F. "A cognitive model of planning." *Cognitive Science Journal,* 1979,275-310.

[7] Hendler, J. "Integrating marker-passing and problem solving" in COGSCI-85.

[8] Kolodner, J., and Simpson, R, "Experience and problem solving: a framework" in COGSCI-84.

[9] Kolodner, J., Simpson, R., and Sycara-Cranski, K, "A process model of case-base reasoning in problem solving" in IJCAI-1985.

[10] Lakoff, G. *Women, Fire, and Dangerous Things.* University of Chicago Press, 1987.

[11] McCarthy, J. and Hayes, P.J. "Some philosophical problems from the standpoint of Artificial Intelligence" in Meltzer and Michie (Eds.) *Machine Intelligence 4.* Edinburgh University Press, 1969.

[12] Newell, A. "Artificial Intelligence and the Concept of Mind" in R. Schank and K. Colby (Eds.) *Computer Models of Thought and Language.* W.H. Freeman and Co., 1973.

[13] Rissland, E. and Ashley, K. "Hypotheticals as heuristic device" in AAAI-86.

[14] Rosch, E. "Prototype Classification and Logical Classification: The Two Systems." Paper presented at a meeting of the Jean Piaget Society, Philadelphia, 1981.

[15] Sacerdoti, E. *A structure for plans and behavior.* Elsevier North-Holland, 1977.

[16] Schank, R.C. "Language and Memory." *Cognitive Science Journal,* 1980,4,243-284.

[17] Smith, E. and Medin, D. *Categories and Concepts.* Havard University Press, 1981.

[18] Stanfill, C. and Waltz, D. "Toward memory-based reasoning." *Communications of the ACM,* 1986,29,1213-1239.

[19] Sussman, G. *A computer model of skill acquisition.* Elsevier North-Holland, 1975.

[20] Tulving, E. "Episodic and Semantic Memory" in E. Tulving and W. Donaldson (Eds.), *Organization of memory.* Academic Press, 1972. Compiler."

[21] Wilensky, R. *Planning and Understanding.* Addison- Wesley, 1983.