Using Plan Reasoning in the Generation of Plan Descriptions *

R. Michael Young
Intelligent Systems Program
University of Pittburgh
Pittsburgh, PA 15260
myoung+@pitt.edu

Abstract

Previous work on the generation of natural language descriptions of complex activities has indicated that the unwieldy amount of text needed to describe complete plans makes for ineffective and unnatural descriptions. We argue here that concise and effective text descriptions of plans can be generated by exploiting a model of the hearer's plan reasoning capabilities. We define a computational model of the hearer's interpretation process that views the interpretation of plan descriptions as refinement search through a space of partial plans. This model takes into account the hearer's plan preferences and the resource limitations on her reasoning capabilities to determine the completed plans she will construct from a given partial description.

Introduction

A number of natural language systems have been developed for the generation of textual descriptions of plans (Mellish & Evans 1989; Vander Linden 1993; Haller 1994). However, systems have been limited in their ability to deal with the large amount of detail found in complex activities: either these systems dealt exclusively with artificial plans of limited size or have generated verbose text describing more realistic plans. The quality of a textual description is strained when that description contains an inappropriate amount of detail. Providing a hearer with too much detail may needlessly cause her to eliminate from consideration compatible alternate plans or may overtax her attentional constraints (Walker 1996). Too little detail, alternatively, may allow the hearer to infer that the speaker is describing a plan that is inconsistent with the speaker's actual plan. Providing too little detail may so underconstrain the interpretation that the hearer's plan reasoning resources are overtaxed. For systems responsible for automatically generating descriptions of plans, understanding the interaction between the quality of the description and the quantity of information it contains is essential

In the approach we describe here, plans are represented by collections of components and the task of describing a plan involves communicating these components. The problem that this research addresses is how to determine an appropriate subset of the components of a plan to communicate as the plan's description. A principal claim we make is that effective description of one plan can be achieved by describing a second *partial* plan that is appropriately related to the first. The partiality of the second plan must be chosen so that the hearer can reconstruct the first from it based on the hearer's knowledge about plans and planning. The hearer must be able to complete the description in much the same way that a planning system completes a partial plan.

Exploiting a Model of Plan Reasoning

This work addresses the communication of plan descriptions in a context we call *plan identification*. In this context, a speaker describes a plan P, called the *source plan*, in order to identify P as a solution to what the speaker believes is a mutually understood planning problem. When identifying a plan, a speaker provides a description of P that is sufficient for the hearer to single out P (or a plan sufficiently close to P) as the indicated plan in the space of possible solutions. The description the speaker provides will contain a description of a subset of the plan components found in P – the speaker constructs this subset by anticipating how the hearer will reconstruct a complete plan from the partial plan defined in the description.

In this paper we will only consider utterances that describe the presence of a component in the source plan.[1] In instructional text, for instance, these types of utterances are often realized as imperatives like "Do action α." The problem we address here is the selection of a subset of the components of a source plan

*Research supported by the Office of Naval Research, Cognitive and Neural Sciences Division, DoD FY92 Augmentation of Awards for Science and Engineering Research Training (ASSERT) Program (Grant Number: N00014-93-I-0812)

[1] This formalism is readily extendable to other types of utterances – see below for a brief discussion.

P sufficient to identify P to the hearer. In our approach, a speaker uses the model of the hearer's plan reasoning capabilities to select a subset with greater or fewer components depending on (at least) two factors. First, the hearer's plan reasoning resources may limit her ability to find completions; the constraints in a description may be so sparse that finding the completions of the constraints is too great a task for the hearer. Candidate subsets that overtax the hearer's abilities can be eliminated from consideration.

The second factor determining the content of a description is the amount of variance the speaker can tolerate between the plan he is describing and the plan the hearer subsequently will form. In general, there may be a number of plans closely related to a source plan which the speaker would be happy for the hearer to identify. The degree to which these closely related plans vary from the actual source plan is dependent upon the measure of *acceptability* that the speaker uses. In the limiting case, acceptability corresponds to identity, although for many domains acceptability may be a much weaker notion. Given the hearer's plan preferences, the particular constraints in a description may guide him to solutions that are unacceptable to the speaker. Subsets that define planning problems where unacceptable plans are likely to be selected can also be eliminated.

We will use a representation of the planning process referred to as *plan-space search* (Kambhampati 1993). Plan-space search provides a flexible representation of partial plans and the types of planning operations hearers may perform when interpreting a partial plan description. In addition, plan-space search characterizes the planning activity of a wide class of current planning systems – developing a text generation system built on plan-space search allows us to apply these techniques to any plan representation that can be characterized in this way. In plan-space search, each node in the search space is defined by a partial plan and denotes a set of plans; this set, called the node's *candidate set*, corresponds to the class of all legal completions of the partial plan at that node. Search through the plan space is performed by refining the partial plan at a given node. Refinements correspond to one of a well-specified set of plan-construction operations (e.g., adding a new step to establish an open precondition). Each refinement of a parent node creates a child node whose candidate set is a subset of the parent's. The plan space forms a graph whose single root node is an empty plan and whose leaf nodes are either inconsistent plans or solutions to the planning problem.

In refinement search, an evaluation function is used to characterize the candidate set of each node encountered in the search, mapping the node to one of three values. When a node evaluates to `FAIL`, there is no plan in the node's candidate set that solves the planning problem. Consequently, the node must be pruned from the search space and the search algorithm must backtrack. When a node evaluates to a complete plan contained in the candidate set, that plan solves the planning problem and the search algorithm can return this plan. When a node evaluates to $\perp$, the evaluation function cannot determine if a solution is contained in the candidate set and further refinement is needed.

In this paper, we will use a partial-order, causal link planner called DPOCL (Young, Pollack, & Moore 1994). This planner extends the UCPOP planner (Penberthy & Weld 1991) by incorporating a hierarchical plan representation. Use of this representation mirrors the hierarchical, incremental development of plans indicated in the manner that people *talk* about planning (Dixon 1987). Because DPOCL is not a system built especially for the generation of task descriptions (i.e., it is a domain-independent planning algorithm), DPOCL plans contain sufficient structure to ensure their executability. Consequently, they serve as strong test cases for the generation of plan descriptions. In addition, DPOCL is readily characterized as a plan-space search algorithm.

A Hearer Model Based on Plan Reasoning

In this section we propose a model to be used for determining an appropriate description of a plan – using this model involves determining specific inferences to be drawn by the hearer from any candidate description. In particular, our model anticipates the plan reasoning that the hearer undertakes to complete the partial description the speaker provides.

Definitions

The computational model we use here is made of a number of components representing the planning algorithms and action representations used by a speaker when modeling the domain of discourse and the hearer's model of the same domain. In our approach, the speaker has a planning model representing his own plan reasoning capabilities and a separate model of the hearer's plan reasoning capabilities. [2]

A planning problem consists of a complete specification of the problem's initial state and the goal state and a complete specification of the domains's action and decompositions operators.

Definition 1 (Planning Problem) *A planning problem $\mathcal{PP}$ is a three-tuple $\prec P_0, \Lambda, \Delta \succ$, where P_0 is a plan specifying only the initial state and the goal state, Λ is the planning problem's set of action operator definitions and Δ is the set of decomposition operator definitions.*

As described above, a plan-space planning algorithm searches a plan space to find a solution to a

[2] We will refer to the speaker's planning model as the *speaker model* and the speaker's model of the hearer's planning model as the *hearer model*.

planning problem. Typical implementations produce a plan graph during this search representing the portion of the plan space searched to that point.

Definition 2 (Plan Graph) *A plan graph $G_A = \prec n, a \succ$ for some planning algorithm $\mathcal{A}$ is a singly-rooted, strongly connected graph with nodes n and arcs a. Each node $n_i \in n$ is a plan defined by algorithm $\mathcal{A}$ and an arc $n_i \rightarrow n_j$ appears in a precisely when n_j is a refinement of n_i using algorithm $\mathcal{A}$.*

During the planning process, the hearer model employs a heuristic evaluation function to direct search through the space of plans. This function ranks plans that appear in the fringe of the current plan graph; search proceeds by expanding those fringe nodes that are ranked most promising.

Definition 3 (Plan Ranking Function) *For any plan p and plan graph G_A, $G_A = \prec n, a \succ$ and $p \in n$, a plan ranking function f maps p and G_A into a set of plans such that 1) f partitions the plans in n into a totally ordered set of sets of plans, 2) this total order has a single minimal element and 3) each plan in n must be assigned by f into precisely one of these sets.* [3]

For ease of reference, we will identify the total ordering on these sets with the non-negative integers; plans that are assigned a lower number are more preferred than plans assigned a higher ranking.

In order to model the resource limits of a hearer when she is interpreting a description, the hearer's limitations will be represented by a *search limit function* that accepts as input a plan graph representing the space already explored during a plan-space search. The function returns T if the plan graph exceeds the hearer's search limit and returns F if it does not.

Definition 4 (Search Limit Function) *A search limit function d maps a plan graph G onto the set $\{T, F\}$. For any agent a, $d_a(G) = T$ precisely when G exceeds the planning resource limit for a.*

The hearer model's planning system combines a particular planning algorithm, a search limit function and an evaluation function.

Definition 5 (Planning System) *A planning system $\mathcal{PS}$ is a three-tuple $\prec \mathcal{A}, d, f \succ$ where $\mathcal{A}$ is a plan-space search algorithm, d is a search limit function and f is a plan evaluation function.*

Finally, a planning model $\mathcal{PM}$ is a pair consisting of a planning problem and a planning system.

Definition 6 (Planning Model) *A planning model $\mathcal{PM}$ consists of a planning problem $\mathcal{PP}$ and a planning system $\mathcal{PS}$.*

[3] The system designer is free to rank plans using any criteria and to compute that ranking using any mechanism. See (Elzer, Chu-Carroll, & Carberry 1994).

Complete plans assigned the rank 0 by an evaluation function in a planning model $\mathcal{PM}$ are called the *preferred plans* in $\mathcal{PM}$. In this work we will use as a measure of acceptability the difference between the value of the speaker's evaluation function f applied to a plan and f's value when applied to the speaker's source plan P, We will assume that the speaker has some value δ that serves as a measure of the amount of variance from P that he will tolerate. Let G_s be the plan graph in which P was found by the speaker, and let G_h be some plan graph constructed by the hearer while solving the same planning problem. The set of *acceptable* plans (or simply the *acceptance set*) for a given source plan P contains precisely those plans P' in G_h such that $|f_s(P, G_s) - f_s(P', G_h)| \leq \delta$.

Constraints on the Hearer Model

There are a number of constraints that must be placed on the plans produced by either the speaker or the hearer model. First, as described earlier, we will use DPOCL to model the planning algorithm of the hearer. That is, $\mathcal{A}_h$ = DPOCL. Second, the planning algorithms of both speaker and hearer will be constrained to produce only complete plans that contain no unnecessary steps. We assume that the definition of the planning problem in the speaker and hearer models are identical. That is, the specification of the initial and goal states are the same. Furthermore, the discussion here will not deal with any incompatibility between the speaker and the hearer models' representations of the operators in our domain. [4]

Putting the Hearer Model to Use

This hearer model is put to use during the selection of the contents of a plan description. The constraints in this description create a new planning problem for the hearer, one in which the empty plan P_0 is replaced as the root node by the partial plan characterized by the description. We will refer to this new plan as $\mathcal{P_D}$. $\mathcal{P_D}$ has the same initial and goal states as P_0 but has some amount of plan detail already filled in. As a result, the characteristics of the plan space below this node differ from that of the plan space of the original problem.

By examining the manner in which this new plan space will be searched in the hearer model, the speaker can determine the efficacy of the corresponding description. To be cooperative, a speaker should select a set of plan constraints $\mathcal{P_D}$ such that

- **Acceptability:** the speaker believes that all completions of $\mathcal{P_D}$ of equal or greater preference to the

[4] In general, there may be considerable variance between the planning model used by the speaker and that assumed to be used by the hearer. Our approach does not commit to a particular policy for reconciling differences between speaker and hearer models but instead allows implementors to impose policy as their applications dictate.

source plan in the hearer model also occur in the acceptance set of the speaker, and

- **Resource Adequacy:** the speaker believes that at least one such acceptable completion exists in the plan-space of the hearer model within the bound d_h.

When identifying a plan, the speaker should describe a partial plan whose completions in the hearer model are acceptable with respect to the source plan. One interpretation of Grice's maxim of quantity (Grice 1975) suggests that the speaker must determine a *minimal* set of constraints that meet these requirements. To find a minimal subset of plan constraints to use as a description, the approach defined here uses the planning system of the hearer model to evaluate candidate descriptions. To determine if a collection of plan constraints describe a set of plans that are all acceptable with respect to the source plan, we can initialize the hearer model's planning problem using a subset of the source plan's components and search the space of plans rooted at that node. To find a description that obeys the maxim of quantity, we begin our search with the empty subset and increase the size of the initial component set until we reach the first set defining a plan space where every preferred plan is acceptable.

A Sample Problem

This section examines three descriptions for the same example planning problem. The planning problem we will use involves travel from London to Paris. [5] In this domain, there are three basic ways to travel from London to Paris: by train, by plane and by automobile. Each option involves the specification of some further detail; one can fly to either of the two airports in Paris, take the train to Paris either by ferrying across the English Channel or traveling directly using the Channel Tunnel (on the *Eurostar*) or drive to Paris, again by taking the ferry across the Channel.

The complete plan space for this planning problem is shown in figure 1. Each node in this graph represents a (possibly partial) plan; the graph is rooted at P_0, the null plan that describes only the initial state (being in London) and the goal state (being in Paris). Each arc between two nodes in the graph indicates a refinement of the plan at the first node to form a new plan at the arc's second node. The leaf nodes in this graph are all solutions to the planning problem and are labeled with text giving a rough indication of their structure. Each node is also labeled with an integer indicating the order that the hearer's plan search function f_h will search the space (described further below).

We will define the hearer model as follows. For illustrative purposes, we will assume a limited resource

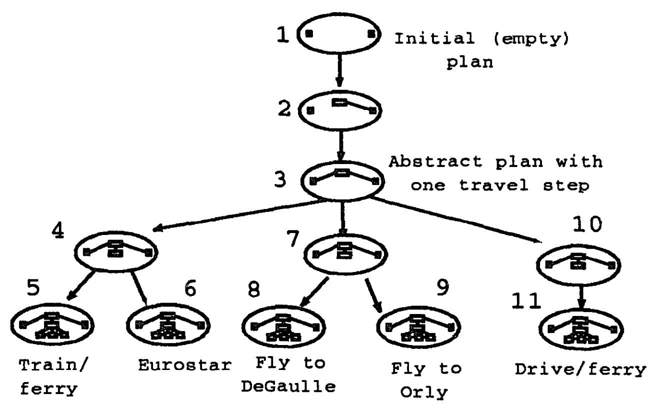

Figure 1: Complete Plan Space for Travel Problem.

bound on the hearer's plan reasoning, bounding the search she can perform to graphs with fewer than 5 nodes. Formally, for any graph $G = \prec n, a \succ$, $d_h(G) = F$ precisely when $|n| < 5$.

The speaker's source plan P is the plan to fly to Paris by taking a plane from Heathrow to De Gaulle (numbered #8 in figure 1). We will assume that the speaker has two simple but strong factors that effect his plan preferences: strong preferences *against* any plan that involves train or ferry travel. The speaker's plan evaluation function f_S and his measure of variance δ will be set such that the only acceptable plans are those numbered #8 and #9 in figure 1.

Providing Too Much Detail Consider the following description:

> **Description 1** `To travel to Paris from London, take the Tube to Heathrow. Next, take a plane to De Gaulle. Then take a bus from De Gaulle to Paris. In order to be in London before taking the Tube, use the effect of starting in London. In order to be at Heathrow before taing the plane to Paris, use the effect of begin at Heathrow after taking the Tube. In order to be at De Gaulle before taking a bus to Paris, use the effect of taking a plane from Heathrow to De Gaulle. In order to be in Paris, use the effect of taking a bus from De Gaulle into Paris.`

This description provides enough detail so that it specifies exactly one completed plan: the source plan P that appears as a leaf node (#8) in the original plan space from figure 1. The new plan space rooted at this node contains just this single plan.

To evaluate this candidate description, the speaker uses the hearer model described above to complete this

[5] In these examples, we will use a simplified version of the DPOCL planner and its plan representation to limit the length of the discussion. The techniques we present here, however, are applicable to planning problems of arbitrary complexity.

plan. Because the plan specified by the description is already complete, no search is required; this plan is the only plan in the set of preferred solutions in the hearer model. The plan is acceptable, since it is the source plan.

By including so much detail in the description, the speaker eliminates other acceptable plans – in this case the plan to fly into Paris's De Gaulle airport. Moving the root node of the new plan space farther from the source plan would make for a more concise description while including other acceptable plans in the hearer model's preference set. However, as we show in the next section, this may also result in the inclusion of unacceptable plans in the preference set.

Providing Insufficient Detail Consider the following description:

Description 2 `In order to be in Paris, travel to Paris from London.`

This description describes a plan containing only the abstract TRAVEL action with no additional detail. This new plan space is rooted at the node labeled #3 in figure 1 With a search limit function constraining the hearer model's plan graph to contain no more than 5 nodes, the hearer model will only find two solutions to the planning problem: nodes #5 and #6. Both of these nodes are unacceptable given the definition of f_S described earlier. There are, then, no plans in the preferred set of the hearer model that are also acceptable to the speaker, making description 2 unacceptable.

Locating an Appropriate Description

Effective descriptions may make reference to any or all of the components of a plan. Consequently, candidate descriptions lie in the power set of the constraints present in the speaker's source plan. The previous examples use descriptions that lie on either side of an effective, concise description for the source plan. One obvious technique for finding the minimal set of constraints that successfully describes the plan is to use a brute force search algorithm: consider each set in the power set of the constraints of the source plan. This technique would be initialized with the initial, null plan $\mathcal{P}_\mathcal{D}$ and would incrementally evaluate sets of constraints, always considering the unexamined sets with smallest cardinality next. The algorithm would halt with it had either found an acceptable plan or exhausted the power set of plan constraints.

Using this technique, it's possible to locate a set of plan constraints corresponding to the following description:

Description 3 `In order to be in Paris, travel to Paris from London. Travel to Paris by flying.`

This description describes a plan that contains the TRAVEL step, and a decomposition for that step involving a FLY action. The plan is partial, since it does not specify which airport to fly to. This new plan space is rooted at node labeled #7 in figure 1. The hearer model will search the plan space below this node and find two solutions to the planning problem: nodes #8 and #9. Both of these nodes are acceptable (in fact, they are the only two), making the description acceptable. Any other description of similar or lesser size would either be rooted at one of node #7's siblings or its parents (with spaces that would either lead to unacceptable solutions in the preference set of the hearer model or that would contain no solutions at all) description 3 is minimal as well.

The technique of selecting minimal descriptions using exhaustive search corresponds to Dale and Reiter's full brevity interpretation of Grice's maxims used in the generation of referring expressions (Dale & Reiter 1995). This approach has two weaknesses. First, As Dale and Reiter point out, the approach is computationally expensive, making it unappealing for describing complicated plans or plan spaces. Second, it is not guaranteed to isolate a unique description. For any given planning problem, there may be a number of acceptable descriptions all of minimal size. This simple technique is unable to distinguish between such competing descriptions.

In these cases, the plan constraints themselves may suggest heuristics for choosing between candidate descriptions. For instance, partial plans that are more *referentially coherent* (Long & Golding 1993; Kintsch 1988), that is, whose plan graphs have fewer strong components, may be preferred for explanation over those that are not. Work in the comprehension of narrative texts (Long & Golding 1993; Graesser, Lang, & Roberts 1991; Graesser, Roberston, & Anderson 1981; Trabasso & van den Brock 1985) describes types of inferences drawn from descriptions of actions. It is possible that these cognitive models can given computational definitions in plan identification.

Related Work

Several researchers interested in task-related discourse have employed action representations based on AI plans. The principal work on explaining plans produced by AI planning systems is described by Mellish and Evans (Mellish & Evans 1989). Their system takes a plan produced by the NONLIN planner (Tate 1977) and produces a textual description of the plan. Their system generates clauses describing each component of the input plan and, as Mellish and Evans point out, this often results in unnatural descriptions containing too much detail.

Other projects (Vander Linden 1993; Delin *et al.* 1994) produce more concise texts describing activities, but rely on simplified models of plans whose size and complexity were limited. Dale's dissertation (Dale 1992), focusing on the generation of anaphoric referring expressions in text describing cooking plans, avoided the generation of overly detailed descriptions by a combination of domain-specific techniques (e.g., linguistic information about the verbs associated with actions) and domain-independent ones (e.g., exploiting focus constraints within the text being produced).

Conclusions and Future Work

This paper has defined a computational model of a hearer's plan reasoning capabilities that is useful for selecting between competing candidate descriptions. By viewing the hearer's interpretation of a partial description as the task of searching for a completion in a space of plans, we have been able to provide a formal account for the requirements of this task. The requirements are described in terms of the hearer's planning algorithm, her individual plan preferences and any resource limits placed on her planning capabilities. This model characterizes a number of domain-independent planning algorithms; as a result, the model can potentially be used to generate descriptions of plans produced from a number of automatic planning systems.

Future work will address a number of issues. We will examine the use of additional forms of constraints in text descriptions beyond those discussed here (e.g., negative imperatives) and their role in bounding the search space that the hearer model must deal with. In addition, our future work will explore ways to extend this model to contexts where groups of agents use dialog to coordinate their plan-related beliefs. Finally, we will investigate techniques for reconciling differences between a speaker's model of the hearer and the methods actually employed by the hearer during interpretation of a partial plan description.

Acknowledgements

The author thanks Johanna Moore, Martha Pollack and the reviewers for their helpful comments.

References

Dale, R., and Reiter, E. 1995. Computational interpretations of the Gricean Maxims in the generation of referring expressions. *Applied Artificial Intelligence Journal* 9. to appear.

Dale, R. 1992. *Generating referring expressions: Constructing descriptions in a domain of objects and processes.* Cambridge, Massachusetts: MIT Press.

Delin, J.; Hartley, A.; Paris, C.; Scott, D.; and Vander Linden, K. 1994. Expressing procedural relationships in multilingual instructions. In *Proceedings of the Seventh International Workshop on Natural Language Generation*, 61–70.

Dixon, P. 1987. The structure of mental plans for following directions. *Journal of Experimental Psychology: Learning, Memory and Cognition* 13:18–26.

Elzer, S.; Chu-Carroll, J.; and Carberry, S. 1994. Recognizing and utilizing user preferences in collaborative consultation dialogues. In *Proceedings of the Fourth International Conference on User Modeling*, 19 – 24.

Graesser, A.; Lang, K.; and Roberts, R. 1991. Question answering in the context of stories. *Journal of Experimental Psychology: General* 120:254–277.

Graesser, A.; Roberston, S.; and Anderson, P. 1981. Incorporating inferences in narrative representations: a study of how and why. *Cognitive Psychology* 13:1–26.

Grice, H. P. 1975. Logic and conversation. In Cole, P., and Morgan, J. L., eds., *Syntax and Semantics III: Speech Acts.* New York, NY: Academic Press. 41–58.

Haller, S. 1994. *Interactive Generation of Plan Descriptions and Justifications.* Ph.D. Dissertation, State University of New York at Buffalo.

Kambhampati, S. 1993. Planning as refinement search: A unified framework for comparative analysis of search space size and performance. Technical Report 93-004, Arizona State University, Department of Computer Science and Engineering.

Kintsch, W. 1988. The role of knowledge in discourse comprehension: a construction-integration model. *Psychological Review* 95:163–182.

Long, D. L., and Golding, J. M. 1993. Superordinate goal inferences: Are they automatically generated during comprehension? *Discourse Processes* 16:55–73.

Mellish, C., and Evans, R. 1989. Natural language generation from plans. *Computational Linguistics* 15(4).

Penberthy, J. S., and Weld, D. 1991. UCPOP: A sound, complete partial order planner for ADL. In *Proceedings of the Third International Conference on Knowledge Representation and Reasoning.*

Tate, A. 1977. Generating project networks. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 888 – 893.

Trabasso, T., and van den Brock, P. 1985. Causal thinking and the representation of narrative events. *Journal of Memory and Language* 24:612–630.

Vander Linden, K. 1993. *Speaking of Actions: Choosing Rhetorical Status and Grammatical Form in Instructional Text Generation.* Ph.D. Dissertation, University of Colorado, Department of Computer Science.

Walker, M. 1996. The effect of resource limits and task complexity on collaborative planning in dialog. *Artificial Intelligence.* to appear.

Young, R. M.; Pollack, M. E.; and Moore, J. D. 1994. Decomposition and causality in partial order planning. In *Proceedings of the Second International Conference on AI and Planning Systems*, 188–193.