A THREE-LEVEL MODEL FOR PLAN EXPLORATION

Lance A. Ramshaw
Department of Computer Science
Bowdoin College
Brunswick, ME 04011
Internet: ramshaw@polar.bowdoin.edu

ABSTRACT

In modeling the structure of task-related discourse using plans, it is important to distinguish between plans that the agent has adopted and is pursuing and those that are only being considered and explored, since the kinds of utterances arising from a particular domain plan and the patterns of reference to domain plans and movement within the plan tree are quite different in the two cases. This paper presents a three-level discourse model that uses separate domain and exploration layers, in addition to a layer of discourse metaplans, allowing these distinct behavior patterns and the plan adoption and reconsideration moves they imply to be recognized and modeled.

DISCOURSE MODEL LEVELS

In task-related discourse, much of the discourse structure derives directly from the task structure, so that a model of the agent's plans can serve as a useful discourse model, with discourse segment boundaries mapping to sub-plan boundaries and the like. This simple model works well in applications like expert-apprentice dialogues, where a novice agent is currently pursuing a single domain plan. Since the discourse tracks the domain plan so closely in such cases, it is fairly easy to make the links between the agent's queries and the relevant domain plans.

But this single-level model is not rich enough to handle phenomena like clarification subdialogues or plan revision, as seen in the work of Litman, Carberry, and others. Litman's model [Lit85, LA87, LA90] employs a stack of discourse metaplans on top of the base domain plan, producing a two-level model that can handle clarification subdialogues and other discourse phenomena that go beyond the domain plans. Carberry [Car90]

adds an independent stack of discourse goals, for similar reasons.

In earlier work [Ram89a], I explored the use of a different kind of metaplan to model what I called the problem-solving level, where the agent is exploring possible plans, rather than pursuing an adopted one. Such plan exploration contexts, which can include comparison between alternative plans or consideration of plans in hypothetical circumstances, are quite different from adopted domain plan execution contexts, both in terms of the reference patterns to domain plans and in terms of the kinds of utterances that are generated.

This paper describes an effort to combine these earlier approaches into a three-level model where the discourse metaplans can be rooted in either the exploration or domain plan levels, so that both kinds of plan-related behavior can be modeled. Such a model can capture the differences between the plan exploration level and the domain level in terms of the appropriate plan recognition and query generation processes, thus broadening the range of discourse phenomena that can be modeled. It also allows us to model shifts between levels, as the agent explores, adopts, and reconsiders particular plans.

THE NATURE
OF PLAN EXPLORATION

Cohen and Levesque [CL90] have recently pointed out the theoretical problems that arise, while using a planning system to model a rational agent, from failing to distinguish between a system's plans and its intentions, since agents frequently form plans that they never adopt. It is this same distinction that motivates the division proposed here of the domain-plan-related portion of the discourse model into separate levels representing first those domain plans and goals which the agent has adopted and is pursuing, here called the domain

layer, and second those which the agent is exploring but has not adopted, the exploration layer. While the same domain plans give structure to these two levels, the resulting discourse phenomena including the space of relevant queries based on those plans and the patterns of references to plans in the plan tree are quite different.

QUERY TYPES

One clear difference can be seen in the content of utterances arising on those different layers from the same underlying domain plan. For example, in a banking context, consider these two queries:

What is the interest rate on your passbook account?

To whom do I make out the check for the initial deposit on a passbook account?

The interest rate query is an example of a query based in an exploration context, since the interest rate on an account affects its desirability compared to other accounts but has no instrumental relevance to any of the plan steps involved in opening an account. The query about the check payee, on the other hand, plays a local instrumental role in the open-account plan, but has no relevance outside of that particular subplan.

Some queries, of course, can arise in either kind of context. For example,

What's the minimum initial deposit for a passbook account?

could be either an exploration level query from an agent weighing the comparative advantages of a passbook account versus a CD, or it could be a domain level query from an agent who had already decided to open a passbook account, and who needed to know how large a check to write to open it. Thus the context model for that query is ambiguous between the two interpretations.

There are also whole classes of queries that may be generated on the plan exploration level but that do not arise when agents are pursuing an adopted domain plan, as when an agent asks queries about the possible plans for a goal or about possible fillers for a variable within a plan. For example,

What does it take to open an account?

asks about the subactions in a plan being explored, and

What kinds of accounts do you offer?

asks for possible fillers of the account-class variable. Such queries imply that the agent does not have any fully-instantiated adopted plan in mind.

DOMAIN PLAN REFERENCES

Another difference between the exploration and domain levels comes in the patterns of references to domain plans. An agent pursuing an adopted domain plan has a single subplan in focus and typically shifts that focus in an orderly way related to the sequential steps in that plan. On the other hand, at the exploration level, the possible patterns of movement are much less constrained, and conflicting alternative plans or multiple hypothetical plans may be explored simultaneously. Exploration metaplans can capture these more complex patterns.

For example, agents frequently generate queries that compare particular features of alternative plans for the same goal. For instance, after asking about the interest rate on passbook accounts, the agent might naturally ask about the rate for CD's. This kind of comparison can be modeled by a compare-by-feature plan exploration metaplan, which represents the close discourse connection between the similar features of the two different plans. Such feature by feature comparison would be hard to capture in a model based directly on the domain plans, since the focus would have to jump back and forth between the two alternatives. At each step, such a model would seem to predict further queries about the current plan as more likely than a jump back to a query about the other plan, while at the exploration level, we can have plan comparison metaplans that capture either a plan by plan or a feature by feature approach.

A different kind of complex domain plan reference can occur in a hypothetical exploration query, where the agent explores plans in a context that includes projected states of affairs that are different from her own current world model. Of course there is a sense in which every exploration level query is hypothetical, since it concerns the preconditions or effects of executing a plan to which the agent is not yet committed, but the issue here concerns hypothetical queries that assume more than the adoption of the single plan being explored. While modeling arbitrary hypotheticals requires more than a planning system, there are cases where the hypothetical element in the agent's question can be expressed by assuming the adoption of some other plan in addition to the one currently explored. For example, for a hypothetical query like

If I put $1000 in a 1-year CD and withdrew it in a month, what would be the

penalty?

it seems that an exploration level metaplan could use the purchase-CD plan to create the hypothetical context in which the query about the withdrawal plan penalty is to be understood.

Thus the domain plan references in exploration utterances often do not correspond closely to the shape of the domain level plan tree. Exploration metaplans can supply alternative structures that better capture the more complex reference patterns involved in examples like feature comparisons or hypothetical queries.

MOVES BETWEEN LEVELS

Finally, distinguishing between domain plan and exploration behavior is important so that the system can recognize when the agent moves from one level to the other. If an agent has been asking evaluative queries and then proceeds to ask a pure domain level query about one of those plan options, the system should recognize that the agent has adopted that particular plan and is now actively pursuing it, rather than continuing to evaluate alternatives. Such an adoption of a particular subplan establishes the expectation that the agent will continue to pursue it, perhaps asking further domain level queries, either until it is completed and focus moves on to a new plan or until a plan blockage or second thoughts on the agent's part trigger a reconsideration move back to the exploration level.

We can see the importance of this distinction in the care taken by human agents to make clear the level at which they are operating. Agents use mood, cue phrases, and direct informing acts to keep their expert advisor informed as to whether they are speaking of adopted plans or of ones only being explored, so that the expert's plan tracking and responses can be fully cooperative.

STRUCTURE OF THE MODEL
THE DOMAIN LEVEL

The base level in the model is always a domain plan structure representing the plans and goals that the agent is believed to have adopted and to be currently pursuing. These plans are organized (adapting a technique from Kautz [Kau85]) into a classification hierarchy based on their effects, so that the subplan children of a class of plans represent alternative possible strategies for achieving the given effects. The plans also include traditional action links outlining their action decomposition in terms of more primitive plans.

There are two classes of predicates in a plan definition: *preconditions* which must be true for the plan to be executed but which can be recursively planned for, and *constraints* (using Litman's word—Carberry calls them "applicability conditions") which cannot be recursively planned for. Each predicate is also classified as to its relevance, where internally relevant predicates are those whose bindings must be known in order to execute the plan and external predicates are those whose bindings are relevant when evaluating the plan from the outside. Thus, using the earlier examples, the payee identity for write-check is only internally relevant, the interest rate for open-savings-account is only externally relevant, and the minimum initial deposit feature is relevant both internally and externally. This heuristic classification of predicates is used to indicate which ones can be expanded at the domain vs. exploration levels.

THE EXPLORATION LEVEL

The basic exploration metaplan is explore-plan itself, which takes an instantiated domain plan as its single parameter. The expected default exploration pattern simply follows the domain plan tree shape, exploring the subplans and actions beneath that plan, using the explore-subplan and explore-subaction metaplans. This default behavior is compiled in by linking each exploration level node to the explore-plan nodes for the subplans and subactions of the domain plan it references. Thus when the system models a move to the exploration level from a given domain plan node, the entire subtree of possible plans and actions beneath that node is also instantiated beneath the initial exploration level node.

The more complex exploration level strategies are encoded as metalevel subplans and subactions of explore-plan. For example, compare-subplans is a subplan of explore-plan, and compare-by-feature is in turn one subplan of compare-subplans. The system works from this library of exploration metaplans to create trees of possible contexts beneath each explore-plan node that model these alternative strategies of plan exploration.

THE DISCOURSE LEVEL

The metaplan structure directly underlying an utterance is always a discourse metaplan, though it may be as simple as an ask-value metaplan (like Litman's identify-parameter) directly based in the current domain plan context. As in Litman's approach, phenomena like clarification subdialogues

```
Discourse:
*(ask-sub-plans (open-savings-account agent1 ?bank))

Exploration:
*(explore-plan (conduct-banking-activity agent1))
  *(explore-plan (open-savings-account agent1 ?bank))

Domain:
*(manage-money agent1)
  *(conduct-banking-activity agent1)
```

Figure 1: *What kinds of savings accounts do you offer?*

can be handled by further layers of discourse meta-plans that introduce additional structure above the domain plan. In this three-level model, these discourse layer metaplans can also be based on exploration level plans.

In testing for a match between a given query and a discourse context like ask-value, the discourse metaplans have access to the set of relevant predicates from the base context. In determining that set, the system uses the appropriate relevance criteria depending on whether the base context is at the domain or exploration level. There are also particular discourse plans, such as the ask-fillers plan, that require that their base context be at the exploration level.

OPERATION OF THE MODEL

For each utterance, the system begins from the previous context(s) and searches for a discourse node (based either on a domain or exploration node) that matches the utterance. In the following example, an initial domain level context is assumed, with the default top-level goal being that deduced from the situation of the agent entering a bank and approaching the receptionist, namely, (conduct-banking-activity ?agent) as a subplan of (manage-money ?agent).

The matching context for the initial query,

What kinds of savings accounts do you offer?

is seen in Figure 1, with asterisks used to mark the current focused path. No match in the assumed context is found to this particular query using discourse metaplans based in domain plans, although one can imagine other contexts in which this query could be a step in pursuing an adopted

plan, as in a journalist compiling a consumer's report on various banks. But using the normal plans for banking customers, this query matches only on the exploration level, where the agent is exploring the plan of opening a savings account.

Note that an exploration level match could also be found by assuming that the move to the exploration level occurs at open-savings-account, suggesting that the agent has adopted not just the plan conduct-banking-activity, but also the more specific plan open-savings-account. The system finds both matches in such cases, but heuristically prefers the one which makes the weakest assumptions about plans and goals adopted by the agent, thus preferring the model where the open-savings-account plan is only being explored.

Suppose the agent continues with the query

What's the interest rate on your passbook account?

This is matched by a discourse plan based on exploring one of the subplans of open-savings-account, which was the previous exploration level context, as seen in Figure 2 at #1. The system also explores the possibility of matching to a discourse plan based in a domain level plan, which would imply the agent's adoption of the plan. However, the interest rate feature has only external relevance, and thus cannot match queries on the domain level. This query does finds a second match as the beginning of a compare-by-feature (at #2), but the heuristics prefer the match that is closer to the previous context, while discouraging the one-legged comparison.

The agent's next query,

And the rate for the investment account?

42

Discourse:
* (ask-value agent ?rate (open-passbook-account ...))

Exploration:
* (explore-plan (conduct-banking-activity agent1))
  * (explore-plan (open-savings-account agent1 ?bank))
    * (explore-plan (open-passbook-account agent1 ?bank))   #1
      (compare-by-feature (open-savings-account agent1 ?bank)
        (compare-feature (open-passbook-account ...) (interest-rate-of ...))   #2

Domain:
* (manage-money agent1)
  * (conduct-banking-activity agent1)

Figure 2: *What's the interest rate for the passbook account?*

Discourse:
* (ask-value agent1 ?rate (interest-rate-of ...))

Exploration:
* (explore-plan (conduct-banking-activity agent1))
  * (explore-plan (open-savings-account agent1 ?bank))
    (explore-plan (open-passbook-account agent1 ?bank))
    (explore-plan (open-investment-account agent1 ?bank))   #1
    * (compare-by-feature (open-savings-account ...)(interest-rate-of ...))
      (compare-feature (open-passbook-account ...)(interest-rate-of ...))
      * (compare-feature (open-investment-account ...)(interest-rate-of ...))   #2

Domain:
* (manage-money agent1)
  * (conduct-banking-activity agent1)

Figure 3: *And the rate for the investment account?*

can also be matched in two different ways, as seen in Figure 3. One way (at #1) is based in an explore-plan for open-investment-account, suggesting that the agent has simply turned from exploring one plan to exploring an alternative one. But this query also matches (at #2) as a second leg of the compare-by-feature subplan of explore-plan, where the query is part of the comparison between the two kinds of savings accounts based on the interest rate offered. Since that serves as a close continuation of the feature comparison interpretation of the previous query, the latter interpretation is preferred.

The following two queries

How big is the initial deposit for the passbook account?

And for the investment account?

can be matched by a sibling compare-by-feature subtree, as seen in Figure 4. This approach is thus able to represent the logical feature-by-feature structure of such a comparison, rather than having to bounce back and forth between explorations of the two subplan trees.

The next query,

OK, who do I see to open a passbook account?

makes a substantial change in the context, as

Discourse:
* (ask-value agent1 ?deposit (init-deposit-of ...))
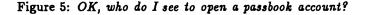

Exploration:
* (explore-plan (conduct-banking-activity agent1))
  * (explore-plan (open-savings-account agent1 ?bank))
    (explore-plan (open-passbook-account agent1 ?bank))
    (explore-plan (open-investment-account agent1 ?bank))
    (compare-by-feature (open-savings-account ...)(interest-rate-of ...))
      (compare-feature (open-passbook-account ...)(interest-rate-of ...))
      (compare-feature (open-investment-account ...)(interest-rate-of ...))
    * (compare-by-feature (open-savings-account ...)(init-deposit-of ...))
      * (compare-feature (open-passbook-account ...)(init-deposit-of ...))
      * (compare-feature (open-investment-account ...)(init-deposit-of ...))

Domain:
* (manage-money agent1)
  * (conduct-banking-activity agent1)

Figure 4: *How big is the initial deposit for the passbook account? And for the investment account?*

Discourse:
* (ask-fillers agent1 ?staff)

Domain:
* (manage-money agent1)
  * (conduct-banking-activity agent1)
    * (open-savings-account agent1 ?bank)
      * (open-passbook-account agent1 ?bank)
        * (fill-out-application agent1 ?staff)

Figure 5: *OK, who do I see to open a passbook account?*

shown in Figure 5. Since the choice of the bank personnel for opening an account is an internal feature that can only be queried on the domain level, the only matches to this query are ones that imply that the agent has adopted the plan that she was previously exploring. Modeling that adoption, the parallel path in the domain tree to the path that was being explored becomes the current domain context, and the matching discourse plan is based there. The cue phrase "OK", of course, is a further signal of this change in level, though not one the system can yet make use of.

In spite of that plan adoption, the agent can later reopen an exploration context concerning a subplan by saying, for example,

What kinds of checks do you have?

She could also raise a query that implies a reconsideration of the previous plan adoption by saying

I forgot to ask whether there are any maintenance charges on this account.

which would reestablish an exploration context of choosing between the passbook and investment accounts.

COMPARISON WITH EXISTING WORK

The general framework of using domain plans to model discourse structure is one that has been

widely pursued and shown to be fruitful for various purposes [All79, AP80, Car84, Car85, GS85, Sid85]. Important extensions have been made more recently in plan classification [Kau85] and in modeling plans in terms of beliefs, so as to be able to handle incorrect plans [Pol86, Pol90].

The most direct precursor of the model presented here is Litman and Allen's work [Lit85, LA87, LA90], which combines a domain plan model with discourse metaplans in a way that can model utterances arising from either the normal flow of domain plans, clarification subdialogues, or cases of domain plan modification. Like exploration metaplans, their discourse plans can handle examples that do not mirror the execution structure of the domain plan. Their system, however, makes the assumption that the agent is pursuing a single domain plan. While the agent can modify a plan, there is no way to capture an agent's exploration of a number of different domain plan possibilities, the use of varying exploration strategies, or the differences between utterances that are based on exploration plans vs. those based on domain plans.

Carberry developed a model [Car90] that is similar to Litman's in combining domain plans with a discourse component, although this model's discourse plans operate on a separate stack rather than as a second layer of the domain plan model. While the mechanisms of her model cover a wide variety of discourse goals, they make no distinction between domain and exploration plans. They are thus also limited to following a single domain plan context at a time.

In earlier work [Ram89a], I presented a model that accounts for the plan refining and query aspects of plan exploration by using a tree of plan-building metaplans, and much of that structure is incorporated in this model. However, that version uses only a single layer of plan-building metaplans, so that it is strictly limited to plan exploration discourse. It thus cannot model queries arising directly from the domain level, nor can it model the moves of plan adoption or reconsideration when the agent switches levels. The plan-building trees in that earlier version are also limited to following the structure of the domain plans, and so are unable to represent comparison by features or other alternative exploration strategies, and that earlier model also lacks a separate discourse component.

Lambert and Carberry [Lam90, LC91] are currently working on a new, three-level approach that has much in common with the one presented here. One interesting difference is that the three levels

in their model form a hierarchy, with discourse plans always rooted in exploration plans. While this may be appropriate for information-seeking discourse, allowing discourse plans to be rooted directly in domain plans can provide a natural way of representing utterances based directly on adopted plans. Overall, their model makes significant contributions on the discourse level, allowing for the recognition of a wide range of discourse plans like expressing surprise or warning. In contrast, the main focus in this work has been on the exploration level, modeling alternative exploration strategies, and plan adoption and reconsideration. It would be fruitful to try to combine the two approaches.

IMPLEMENTATION AND FUTURE WORK

The model presented here has been implemented in a system called Pragma (redone from the earlier Pragma system [Ram89b]) which handles the examples covered in the paper. Since the focus is on modeling plan exploration strategies, the initial context is directly input in the form of a domain plan with its parameter values, and the queries are input as meaning representations. The output after each query is the updated set of context models.

The system has been exercised in the banking and course registration domains, though it is only populated with enough domain plans to serve as a testbed for the plan exploration strategies. The exploration level is the most developed, including metaplans for constraining or instantiating plan variables and for exploring or comparing subplans using various strategies. The discourse level currently includes only the metaplans ask-value, ask-plans, and ask-fillers.

Important next steps include expanding the collection of exploration level metaplans from the samples worked out so far to better characterize the full range of plan exploration strategies that people actually use, validating that collection against real data. It would be particularly interesting to add coverage for the hypothetical queries discussed above, where the assumed event is another known domain plan. The coverage of discourse level metaplans should be expanded, to better explore their interaction with exploration plans. The system should also be made sensitive to other indicators for recognizing moves between the exploration and domain levels besides the class of predicate queried, including verb

mood, cue phrases, and direct inform statements by the agent.

CONCLUSIONS

This work suggests that plan exploration metaplans can be a useful and domain independent way of expanding the range of discourse phenomena that can be captured based on a model of the agent's domain plans. While the more complex exploration strategies complicate the plan recognition task of connecting discourse phenomena with the underlying domain plans, exploration metaplans can successfully model those strategies and also allow us to recognize the moves of plan exploration, adoption, and reconsideration.

References

[All79]   James F. Allen. *A Plan-Based Approach to Speech-Act Recognition.* PhD thesis, University of Toronto, 1979.

[AP80]   James F. Allen and C. Raymond Perrault. Analyzing intention in utterances. *Artificial Intelligence,* 15:143–178, 1980.

[Car84]   Sandra Carberry. Understanding pragmatically ill-formed input. In *Proceedings of the 10th International Conference on Computational Linguistics,* pages 200–206, 1984.

[Car85]   Sandra Carberry. A pragmatics-based approach to understanding intersentential ellipsis. In *Proceedings of the 23rd Annual Meeting of the ACL,* pages 188–197, 1985.

[Car90]   Sandra Carberry. *Plan Recognition in Natural Language Dialogue.* The MIT Press, 1990.

[CL90]   Philip R. Cohen and Hector J. Levesque. Intention is choice with commitment. *Artificial Intelligence,* 42:213–261, 1990.

[GS85]   Barbara Grosz and Candace Sidner. The structures of discourse structure. Technical Report 6097, Bolt Beranek and Newman, 1985.

[Kau85]   Henry A. Kautz. Toward a theory of plan recognition. Technical Report 162, University of Rochester, 1985.

[LA87]   Dianne Litman and James Allen. A plan recognition model for subdialogues in conversation. *Cognitive Science,* 11:163–200, 1987.

[LA90]   Dianne Litman and James Allen. Discourse processing and commonsense plans. In P. Cohen, J. Morgan, and M. Pollack, editors, *Intentions in Communication,* pages 365–388. MIT Press, 1990.

[Lam90]   Lynn Lambert. A plan-based model of discourse understanding. Technical Report 91-04, University of Delaware, 1990.

[LC91]   Lynn Lambert and Sandra Carberry. A tripartite, plan-based model of dialogue. In *Proceedings of the 29th Annual Meeting of the ACL,* 1991.

[Lit85]   Dianne J. Litman. *Plan Recognition and Discourse Analysis: An Integrated Approach for Understanding Dialogues.* PhD thesis, University of Rochester, 1985.

[Pol86]   Martha E. Pollack. *Inferring Domain Plans in Question-Answering.* PhD thesis, University of Pennsylvania, 1986.

[Pol90]   Martha E. Pollack. Plans as complex mental attitudes. In P. Cohen, J. Morgan, and M. Pollack, editors, *Intentions in Communication,* pages 77–103. MIT Press, 1990.

[Ram89a]   Lance A. Ramshaw. A metaplan model for problem-solving discourse. In *Proceedings of the Fourth Conference of the European Chapter of the ACL,* pages 35–42, 1989.

[Ram89b]   Lance A. Ramshaw. *Pragmatic Knowledge for Resolving Ill-Formedness.* PhD thesis, University of Delaware, 1989.

[Sid85]   Candace Sidner. Plan parsing for intended response recognition in discourse. *Computational Intelligence,* 1:1–10, 1985.