

A TRIPARTITE PLAN-BASED MODEL OF DIALOGUE

Lynn Lambert
Sandra Carberry
Department of Computer and Information Sciences
University of Delaware
Newark, Delaware 19716, USA

Abstract¹

This paper presents a tripartite model of dialogue in which three different kinds of actions are modeled: domain actions, problem-solving actions, and discourse or communicative actions. We contend that our process model provides a more finely differentiated representation of user intentions than previous models; enables the incremental recognition of communicative actions that cannot be recognized from a single utterance alone; and accounts for implicit acceptance of a communicated proposition.

1 Introduction

This paper presents a tripartite model of dialogue in which intentions are modeled on three levels: the *domain level* (with domain goals such as *traveling by train*), the *problem-solving level* (with plan-construction goals such as *instantiating a parameter in a plan*), and the *discourse level* (with communicative goals such as *expressing surprise*). Our process model has three major advantages over previous approaches: 1) it provides a better representation of user intentions than previous models and allows the nuances of different kinds of goals and processing to be captured at each level; 2) it enables the incremental recognition of communicative goals that cannot be recognized from a single utterance alone; and 3) it differentiates between illocutionary effects and desired perlocutionary effects, and thus can account for the failure of an inform act to change a hearer's beliefs[Per90]².

2 Limitations of Current Models of Discourse

A number of researchers have contended that a coherent discourse consists of segments that are related to one another through some type of structuring relation[Gri75, MT83] or have used rhetorical relations to generate coherent text[Hov88, MP90]. In addition, some researchers

have modeled discourse based on the semantic relationship of individual clauses[Pol86a] or groups of clauses[Rei78]. But all of the above fail to capture the goal-oriented nature of discourse. Grosz and Sidner[GS86] argue that recognizing the structural relationships among the intentions underlying a discourse is necessary to identify discourse structure, but they do not provide the details of a computational mechanism for recognizing these relationships.

To account for the goal-oriented nature of discourse, many researchers have adopted the planning/plan-recognition paradigm[AP80, PA80] in which utterances are viewed as part of a plan for accomplishing a goal and understanding consists of recognizing this plan. The most well-developed plan-based model of discourse is that of Litman and Allen[LA87]. However, their discourse plans conflate problem-solving actions and communicative actions. For example, their Correct-Plan has the flavor of a problem-solving plan that one would pursue in attempting to construct another plan, whereas their Identify-Parameter takes on some of the characteristics of a communicative plan that one would pursue when conveying information. More significantly, their model cannot capture the relationship among several utterances that are all part of the same higher-level discourse plan if that plan cannot be recognized and added to their plan stack based on analysis of the first utterance alone. Thus, if more than one utterance is necessary to recognize a discourse goal (as is often the case, for example, with warnings), Litman and Allen's model will not be able to identify the discourse goal pursued by the two utterances together or what role the first utterance plays with respect to the second. Consider, for example, the following pair of utterances:

- (1) *The city of xxx is considering filing for bankruptcy.*
- (2) *One of your mutual funds owns xxx bonds.*

Although neither of the two utterances alone constitutes a warning, a natural language system must be able to recognize the warning from the set of two utterances together.

Our tripartite model of dialogue overcomes these limitations. It differentiates among domain, problem-solving, and communicative actions yet models the relationships among them, and enables

¹This material is based upon work supported by the National Science Foundation under Grant No. IRI-8909332. The Government has certain rights in this material.

²We would like to thank Kathy McCoy for her comments on various drafts of this paper.

the recognition of communicative actions that take more than one utterance to achieve but which cannot be recognized from the first utterance alone.

In the remainder of this paper, we will present our tripartite model, motivating why our model recognizes three different kinds of goals, describing our dialogue model and how it is built incrementally as a discourse proceeds, and illustrating this plan inference process with a sample dialogue. Finally, we will outline our current research on modeling negotiation dialogues and recognizing discourse acts such as expressing surprise.

3 A Tripartite Model

3.1 Kinds of Goals and Plans

Our plan recognition framework recognizes three different kinds of goals: domain, problem-solving, and discourse. In an information-seeking or expert-consultation dialogue, one participant is seeking information and advice about how to construct a plan for achieving some *domain goal*. A *problem-solving goal* is a metagoal that is pursued in order to construct a domain plan [Wil81, LA87, Ram89]. For example, if an agent has a goal of earning an undergraduate degree, the agent might have the problem-solving goal of selecting the instantiation of the degree parameter as BA or BS and then the problem-solving goal of building a sub-plan for satisfying the requirements for that degree. A number of researchers have demonstrated the importance of modeling domain and problem-solving goals [PA80, Wil81, LA87, vBC86, Car87, Ram89].

Intuitively, a *discourse goal* is the communicative goal that a speaker has in making an utterance [Car89], such as obtaining information or expressing surprise. Recognition of discourse goals provides expectations for subsequent utterances and suggests how these utterances should be interpreted. For example, the first two utterances in the following exchange establish the expectation that S1 will either accept S2's response, or that S1 will pursue utterances directed toward understanding and accepting it [Car89]. Consequently, S1's second utterance should be recognized as expressing surprise at S2's statement.

S1: *When does CS400 meet?*

S2: *CS400 meets on Monday from 7-9p.m.*

S1: *CS400 meets at night?*

A robust natural language system must recognize discourse goals and the beliefs underlying them in order to respond appropriately.

The plan library for our process model contains the system's knowledge of goals, actions, and plans. Although domain plans are not mutually known by the participants [Pol86b], how to communicate and how to solve problems are common skills that people use in a wide variety of contexts, so the system can assume that knowledge about discourse and problem-solving plans is shared knowl-

edge. Our representation of a plan includes a header giving the name of the plan and the action it accomplishes, preconditions, applicability conditions, constraints, a body, effects, and goals. Applicability conditions represent conditions that must be satisfied for the plan to be reasonable to pursue in the given situation whereas constraints limit the allowable instantiation of variables in each of the components of a plan [LA87, Car87]. Especially in the case of discourse plans, the goals and effects are likely to be different. This allows us to differentiate between illocutionary and perlocutionary effects and capture the notion that one can, for example, perform an inform act without the hearer adopting the communicated proposition.³ Figure 1 presents three discourse plans and one problem-solving and domain plan.

3.2 Structure of the Model

Agents use utterances to perform communicative acts, such as informing or asking a question. These discourse actions can in turn be part of performing other discourse actions; for example, providing background data can be part of asking a question. Discourse actions can take more than one utterance to complete; asking for information requires that a speaker request the information and believe that the request is acceptable (i.e., that the speaker say enough to ensure that the speaker believes that the request is understandable, justified, and the necessary background information is known by the respondent). Thus, actions at the discourse level form a tree structure in which each node represents a communicative action that a participant is performing and the children of a node represent communicative actions pursued in order to perform the parent action.

Information needed for problem-solving actions is obtained through discourse actions, so discourse actions can be executed in order to perform problem-solving actions as well as being part of other discourse actions. Similarly, domain plans are constructed through problem-solving actions, so problem-solving actions can be executed in order to eventually perform domain actions as well as being part of plans for other problem-solving actions.

Therefore, our Dialogue Model (DM) contains three levels of tree structures,⁴ one for each kind of action (discourse, problem-solving, and domain) with links among the actions on different levels. At the lowest level the discourse actions are represented; these actions may contribute to the problem-solving actions at the middle level which,

³Consider, for example, someone saying "I informed you of X but you wouldn't believe me."

⁴The DM is really a mental model of intentions [Pol86b]. The structures shown in our figures implicitly capture a number of intentions that are attributed to the participants, such as the intention that the hearer recognize that the speaker believes the applicability conditions for the just initiated discourse actions are satisfied and the intention that the participants follow through with the subactions that are part of plans for actions in the DM.

Domain Plan-D1: {*_agent earns a minor in _subj*}

Action: Get-Minor(*_agent*, *_subj*)
Prec: have-plan(*_agent*, Plan-D1, Get-minor(*_agent*, *_subj*))
Body: 1. Complete-Form(*_agent*, change-of-major-form, add-minor)
2. Take-Required-Courses(*_agent*, *_subj*)
Effects: have-minor(*_agent*, *_subj*)
Goal: have-minor(*_agent*, *_subj*)

Problem-solving Plan-P1: {*_agent1 and _agent2 build a plan for _agent1 to do _action*}

Action: Build-Plan(*_agent1*, *_agent2*, *_action*)
AppCond: want(*_agent1*, *_action*)
Constr: plan-for(*_plan*, *_action*)
action-in-plan-for(*_laction*, *_action*)
Prec: selected(*_agent1*, *_action*, *_plan*)
know(*_agent2*, want(*_agent1*, *_action*))
knowref(*_agent1*, *_prop*, prec-of(*_prop*, *_plan*))
knowref(*_agent2*, *_prop*, prec-of(*_prop*, *_plan*))
knowref(*_agent1*, *_laction*, need-do(*_agent1*, *_laction*, *_action*))
knowref(*_agent2*, *_laction*, need-do(*_agent1*, *_laction*, *_action*))
Body: 1. for all actions *_laction* in *_plan*, Instantiate-Vars(*_agent1*, *_agent2*, *_laction*)
2. for all actions *_laction* in *_plan*, Build-Plan(*_agent1*, *_agent2*, *_laction*)
Effects: have-plan(*_agent1*, *_plan*, *_action*)
Goal: have-plan(*_agent1*, *_plan*, *_action*)

Discourse Plan-C1: {*_agent1 asks _agent2 for the values of _term for which _prop is true*}

Action: Ask-Ref(*_agent1*, *_agent2*, *_term*, *_prop*)
AppCond: want(*_agent1*, knowref(*_agent1*, *_term*, believe(*_agent2*, *_prop*)))
¬knowref(*_agent1*, *_term*, believe(*_agent2*, *_prop*))
Constr: term-in(*_term*, *_prop*)
Body: Request(*_agent1*, *_agent2*, Informref(*_agent2*, *_agent1*, *_term*, *_prop*))
Make-Question-Acceptable(*_agent1*, *_agent2*, *_prop*)
Effects: believe(*_agent2*, want(*_agent1*, Informref(*_agent2*, *_agent1*, *_term*, *_prop*)))
Goal: want(*_agent2*, Answer-Ref(*_agent2*, *_agent1*, *_term*, *_prop*))

Discourse Plan-C2: {*_agent1 informs _agent2 of _prop*}

Action: Inform(*_agent1*, *_agent2*, *_prop*)
AppCond: believe(*_agent1*, know(*_agent1*, *_prop*))
¬believe(*_agent1*, believe(*_agent2*, *_prop*))
Body: Tell(*_agent1*, *_agent2*, *_prop*)
Make-Prop-Believable(*_agent1*, *_agent2*, *_prop*)
Effects: believe(*_agent2*, want(*_agent1*, believe(*_agent2*, *_prop*)))
Goal: know(*_agent2*, *_prop*)

Discourse Plan-C3: {*_agent1 tells _prop to _agent2*}

Action: Tell(*_agent1*, *_agent2*, *_prop*)
AppCond: believe(*_agent1*, *_prop*)
¬believe(*_agent1*, believe(*_agent2*, believe(*_agent1*, *_prop*)))
Body: Surface-Inform(*_agent1*, *_agent2*, *_prop*)
Make-Statement-Understood(*_agent1*, *_agent2*, *_prop*)
Effects: told-about(*_agent2*, *_prop*)
Goal: believe(*_agent2*, believe(*_agent1*, *_prop*))

Figure 1: Sample Plans from the Plan Library

in turn, may contribute to the domain actions at the highest level (see Figure 3). The planning agent is the agent of all actions at the domain level, since the plan being constructed is for his subsequent execution. Since we are assuming a cooperative dialogue in which the two participants are working together to construct a domain plan, both participants are joint agents of actions at the problem-solving level. Both participants make utterances and thus either participant may be the agent of an action at the discourse level.

For example, a DM derived from two utterances is shown in Figure 3; its construction is described in Section 3.3. The DM in Figure 3 indicates that the inform and the request were both part of a plan for asking for information; the inform provided background data enabling the information request to be accepted by the hearer. Furthermore, the actions at the discourse level were pursued in order to perform a *Build-Plan* action at the problem-solving level, and this problem-solving action is being performed in order to eventually perform the domain action of getting a math minor. The current focus of attention on each level is marked with an asterisk.

3.3 Building the Dialogue Model

Our process model uses plan inference rules[AP80, Car87], constraint satisfaction[LA87], focusing heuristics[Car87], and features of the new utterance to identify the relationship between the utterance and the existing dialogue model. The plan inference rules take as input a hypothesized action A_i and suggest other actions (either at the same level in the DM or at the immediately higher level) that might be the agent's motivation for A_i .

The focusing heuristics order according to coherence the ways in which the DM might be expanded on each of the three levels to incorporate the actions motivating a new utterance. Our focusing heuristics at the discourse level are:

1. Expand the plan for an ancestor of the currently focused action in the existing DM so that it includes the new utterance, preferring to expand ancestors closest to the currently focused action. This accounts for new utterances that continue discourse acts already in the DM.
2. Enter a new discourse action whose plan can be expanded to include both the existing discourse level of the DM and the new utterance. This accounts for situations in which actions at the discourse level of the previous DM are part of a plan for another discourse act that had not yet been conveyed.
3. Begin a new tree structure at the discourse level. This accounts for initiation of new discourse plans unrelated to those already in the DM.

The focusing heuristics, however, are not

identical for all three levels. Although it is not possible to expand the plan for the focused action on the discourse level since it will always be a surface speech act, continuing the plan for the currently focused action or expanding it to include a new action are the most coherent expectations on the problem-solving and domain levels. This is because the agents are most expected to continue with the problem-solving and domain plans on which their attention is currently centered. In addition, since actions at the discourse and problem-solving levels are currently being executed, they cannot be returned to (although a similar action can be initiated anew and entered into the model). However, since actions at the domain level are part of a plan that is being constructed for future execution, a domain subplan already completely developed may be returned to for revision. Although such a shift in attention back to a previously considered subplan is not one of the strongest expectations, it is still possible at the domain level. Furthermore, new and unrelated discourse plans will often be pursued during the course of a conversation whereas it is unlikely that several different domain plans (each representing a topic shift) will be investigated. Thus, on the domain level, a return to a previously considered domain subplan is preferred over a shift to a new domain plan that is unrelated to any already in the DM.

In addition to different focusing heuristics and different agents at each level, our tripartite model enables us to capture different rules regarding plan retention. A continually growing dialogue structure does not seem to reflect the information retained by humans. We contend that the domain plan that is incrementally fleshed out and built at the highest level should be maintained throughout the dialogue, since it provides knowledge about the agent's intended domain actions that will be useful in providing cooperative advice. However, problem-solving and discourse actions need not be retained indefinitely. If a problem-solving or discourse action has not yet completed execution, then its immediate children should be retained in the DM, since they indicate what has been done as part of performing that as yet uncompleted action; its other descendants can be discarded since the parent actions that motivated them are finished. (For illustration purposes, all actions have been retained in Figure 3.)

We have expanded on Litman and Allen's notion of constraint satisfaction[LA87] and Allen and Perrault's use of beliefs[AP80]. Our applicability conditions contain beliefs by the agent of the plan, and our recognition algorithm requires that the system be able to plausibly ascribe these beliefs in recognizing the plan. The algorithm is given the semantic representation of an utterance. Then plan inference rules are used to infer actions that might motivate the utterance; the belief ascription process during constraint satisfaction determines whether it is reasonable to ascribe the requisite beliefs to the agent of the action and, if not, the inference is rejected. The focusing heuristics allow expecta-

tions derived from the existing dialogue context to guide the recognition process by preferring those inferences that can eventually lead to the most expected expansions of the existing dialogue model.

In [Car89] we claimed that a cooperative participant must explicitly or implicitly accept a response or pursue discourse goals directed toward being able to accept the response. Thus our model treats failure to initiate a negotiation dialogue as implicit acceptance of the proposition conveyed by the response. Consider, for example, the following dialogue:

- S1: *Who is teaching CS360 next semester?*
 S2: *Dr. Baker.*
 S1: *What time does it meet?*

Since S1's second utterance cannot be interpreted as initiating a negotiation dialogue, S1 has implicitly accepted the proposition that Dr. Baker is teaching CS360 next semester as true. This notion of implicit acceptance is similar to a restricted form of Perrault's default reasoning about the effects of an inform act [Per90] and is explained further in [Lam91].

3.4 An Example

As an example of how our process model assimilates utterances and can incrementally recognize a discourse action that cannot be recognized from a single utterance, consider the following:

- S1: (a) *I want a math minor.*
 (b) *What should I do?*

A few of the plans needed to handle this example are shown in Figure 1; these plans assume a cooperative dialogue. From the surface inform, plan inference rules suggest that S1 is executing a *Tell* action and that this *Tell* is part of an *Inform* action (the applicability conditions for both actions can be plausibly ascribed to S1) and these are entered into the discourse level of the DM. No further inferences on this level are possible since the *Inform* can be part of several discourse plans and there is no existing dialogue context that suggests which of these S1 might be pursuing. The system infers that S1 wants the goal of the *Inform* action, namely $\text{know}(S2, \text{want}(S1, \text{Get-Minor}(S1, \text{Math})))$. Since this proposition is a precondition for building a plan for getting a math minor, the system infers that S1 wants $\text{Build-Plan}(S1, S2, \text{Get-Minor}(S1, \text{math}))$ and this *Build-Plan* action is entered into the problem-solving level of the DM. From this, the system infers that S1 wants the goal of that action; since this result is the precondition for getting a math minor, the system infers that S1 wants to get a math minor and this domain action is entered into the domain level of the DM. The resulting discourse model, with links between the actions at different levels and the current focus of attention on each level marked with an asterisk, is shown in Figure 2.

The semantic representation of (b) is

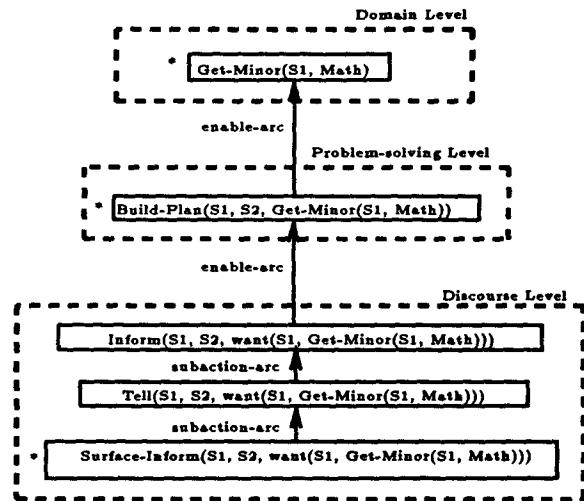


Figure 2: Dialogue Model from the first utterance

Surface-Request(S1, S2, Informref(S2, S1, $_action1$, need-do(S1, $_action1$, $_action2$)))

From this utterance we can infer that S1 is performing a *Request* and thus may be performing an *Ask-Ref* action (since *Request* is part of the body of the plan for *Ask-Ref*) and that S1 may thus be performing an *Obtain-Info-Ref* action (since *Ask-Ref* is part of the body of the plan for *Obtain-Info-Ref*) and that S1 wants the goal of the *Obtain-Info-Ref* action (namely, that S1 know the subactions that he needs to do in order to perform $_action2$), which is in turn a precondition for building a plan. This produces the inference that S1 wants $\text{Build-Plan}(S1, S2, _action2)$ which is an action at the problem-solving level.

The focusing heuristics suggest that the most coherent expectation at the discourse level is that S1's discourse level actions are part of a plan for performing the *Tell* action that is the parent of the action that was previously marked as the current focus of attention in the discourse model. However, no line of inference from the second utterance represents an expansion of this plan. (This means that the proposition was understood without any clarification.⁵) Similarly, no expansion of the plan for the *Inform* action (the other ancestor of the focus of attention in the existing DM) succeeds in linking the new utterance to the DM. (This means that the communicated proposition was accepted without any squaring away of beliefs [Jos82].)

Since the first focusing heuristic was unsuccessful in finding a relationship between the new utterance and the existing dialogue model, the sec-

⁵We are assuming that the hearer has an opportunity to intervene after an utterance. This is a simplification and must eventually be removed to capture a hearer's saving his requests for clarification and negotiation of beliefs until the end of the speaker's complete turn.

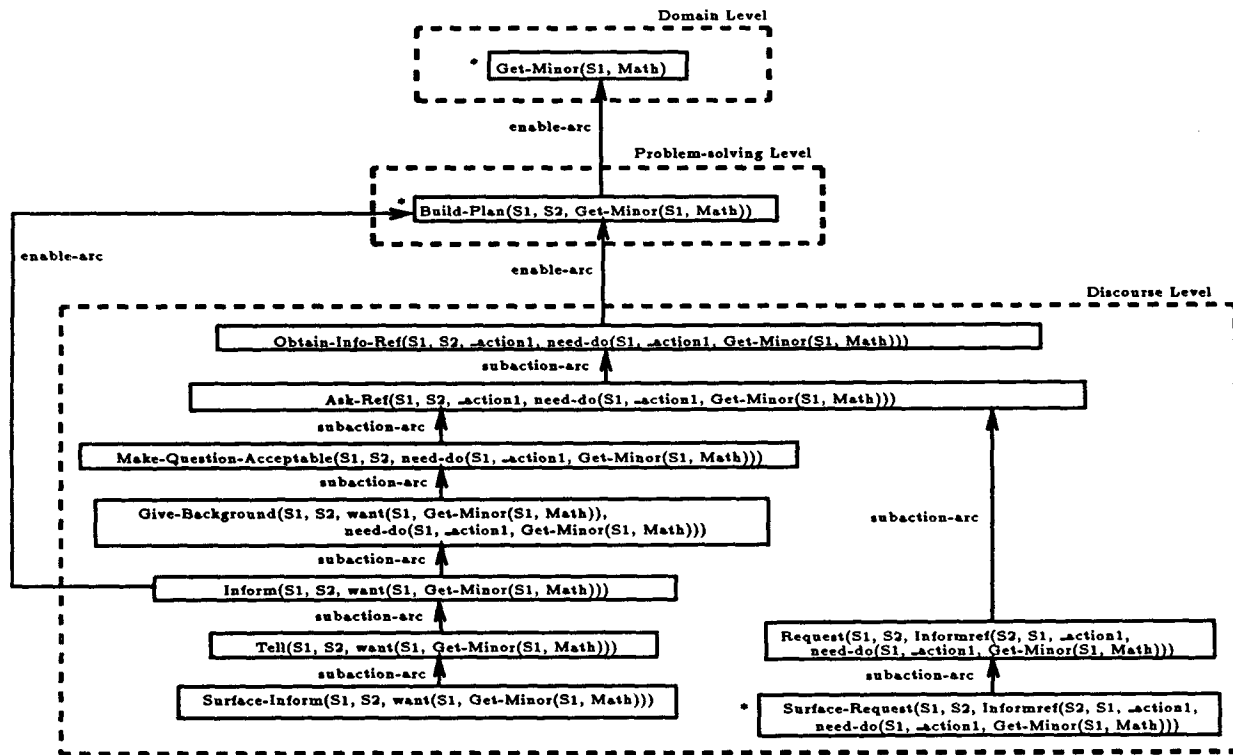


Figure 3: Dialogue Model derived from two utterances

ond focusing heuristic is tried. It suggests that the new utterance and the actions at the discourse level in the existing DM might both be part of an expanded plan for some other discourse action. The inferences described above lead from (b) to the discourse action *Ask-Ref* whose plan can be expanded as shown in Figure 3 to include, as background for the *Ask-Ref*, the *Inform* and the *Tell* actions that were entered into the DM from (a).⁶ The focusing heuristics suggest that the most coherent continuation at the problem-solving level is that the new utterance is continuing the *Build-Plan* that was previously marked as the current focus of attention at that level. This is possible by instantiating *_action2* with *Get-minor(S1, math)*. Thus the DM is expanded as shown in Figure 3 with the new focus of attention on each level marked with an asterisk. Note that S1's overall goal of obtaining information was not conveyed by (a) alone; consequently, only after both utterances were coherently related could it be determined that (a) was part of an overall discourse plan to obtain information and that (a) was intended to provide background data for the request being made in (b).⁷

⁶Note that the actions in the body of *Ask-Ref* are not ordered; an agent can provide clarification and background information before or after asking a question.

⁷An inform action could also be used for other purposes, including justifying a question and merely conveying information.

Further queries would lead to more elaborate tree structures on the problem-solving and domain levels. For example, suppose that S1 is told that Math 210 is a required course for a math minor. Then a subsequent query such as *Who is teaching Math 210 next semester?* would be performing a discourse act of obtaining information in order to perform a problem-solving action of instantiating a parameter in a *Learn-Material* domain action. Since learning the material from one of the teachers of a course is part of a domain plan for taking a course and since instantiating the parameters in actions in the body of domain plans is part of building the domain plan, further inferences would indicate that this *Instantiate-Vars* problem-solving action is being executed in order to perform the problem-solving action of building a plan for the domain action of taking Math 210 in order to build a plan to get a math minor. Consequently, the domain and problem-solving levels would be expanded so that each contained several plans, with appropriate links between the levels.

4 Current and Future Work

We are currently examining the applications that this model has in modeling negotiation dialogues and discourse acts such as convince, warn, and express surprise. To extend our notion of implicit acceptance of a proposition to negotiation di-

alogues, we are exploring treating a discourse plan as having successfully achieved its goal if it is plausible that all of its subacts have achieved their goals and all of its applicability conditions (except those negated by the goal) are still true after the subacts have been executed.

Especially in negotiation dialogues, a system must account for the fact that a user may change his mind during a conversation. But often people only slightly modify their beliefs. For example, the system might inform the user of some proposition about which the user previously held no beliefs. In that case, if the user has no reason to disbelieve the proposition, the user may adopt that proposition as one of his own beliefs. However, if the user disbelieved the proposition before the system performed the inform, then the user might change from disbelief to neither belief nor disbelief; a robust model of understanding must be able to handle a response that expresses doubt or even disbelief at a previous utterance, especially in modeling arguments and negotiation dialogues. Thus, a system should be able to (1) represent levels of belief, (2) recognize how a speaker's utterance conveys these different levels of belief, (3) use these levels of belief in recognizing discourse plans, and (4) use previous context and a user's responses to model changing beliefs.

We are investigating the use of a multi-level belief model to represent the strength of an agent's beliefs and are studying how the form of an utterance and certain clue words contribute to conveying these beliefs. Consider, for example, the following two utterances:

- (1) *Is Dr. Smith teaching CS310?*
- (2) *Isn't Dr. Smith teaching CS310?*

A simple yes-no question as in utterance (1) suggests only that the speaker doesn't know whether Dr. Smith teaches CS310 whereas the form of the question in utterance (2) suggests that the speaker has a relatively strong belief that Dr. Smith teaches CS310 but is uncertain of this. These beliefs conveyed by the surface speech act must be taken into account during the plan recognition process. Thus our plan recognition algorithm will first use the effects of the surface speech act to suggest augmentations to the belief model. These augmentations will then be taken into account in deciding whether requisite beliefs for potential discourse acts can be plausibly ascribed to the speaker and will enable us to identify such discourse actions as expressing surprise. [Lam91] further discusses the use of a multi-level belief model and its contribution in modeling dialogue.

5 Related Work

Ramshaw[Ram91] has developed a model of discourse that contains a domain execution level, an exploration level, and a discourse level. In his model, discourse plans can refer either to the exploration level (corresponding to queries about possible ways of achieving a goal) or to the domain exe-

cution level (corresponding to queries after commitment has been made to achieve a goal in a particular way). In our tripartite model, discourse, problem-solving, and domain plans form a hierarchy with links between adjacent levels. Whereas Ramshaw's exploration level captures the consideration of alternative plans, our intermediate level captures the notion of problem-solving and plan-construction, whether or not there has been a commitment to a particular way of achieving a domain goal. Thus a query such as *To whom do I make out the check?* would be recognized as a query against the domain execution level in Ramshaw's model (since it is a query made after commitment to a plan such as opening a passbook savings account[Ram91]), but our model would treat it as a discourse plan that is executed to further the problem-solving plan of instantiating a parameter in an action in a domain plan — i.e., our model would view the agent as asking a question in order to further the construction of his partially constructed domain plan.

Our tripartite model offers several advantages. Ramshaw's model assumes that the top-level domain plan is given at the outset of the dialogue and then his model expands that plan to accommodate user queries. Our model, on the other hand, builds the DM incrementally at each level as the dialogue progresses; it therefore can handle bottom-up dialogues[Car87] in which the user's overall top-level goal is not explicitly known at the outset and can recognize discourse actions that cannot be identified from a single utterance. In addition, our domain, problem-solving, and discourse plans are all recognized incrementally using basically the same plan recognition algorithm on each level[Wil81]. Consequently, we foresee being able to extend our model to include additional pairs of problem-solving and discourse levels whose domain level contains an existing problem-solving or discourse plan; this will enable us to handle utterances such as *What should we work on next?* (query trying to further construction of a problem-solving plan) and *Do you have information about ...?* (query trying to further construction of a discourse plan to obtain information).

Ramshaw's plan exploration strategies, his differentiation between exploration and commitment, and his heuristics for recognizing adoption of a plan are very important. While our work has not yet addressed these issues, we believe that they are consistent with our model and are best addressed at our problem-solving level by adding new problem-solving metaplans. Such an incorporation will have several advantages, including the ability to handle utterances such as

If I decide to get a BA degree, then I'll take French to meet the foreign language requirement.

In the above case, the speaker is still exploring a plan for getting a BA degree, but has committed to taking French to satisfy the foreign language requirement should the plan for the BA degree be adopted. It does not appear that Ramshaw's model

can handle such contingent commitment. This enrichment of our problem-solving level may necessitate changes to our focusing heuristics.

6 Conclusions

We have presented a tripartite model of dialogue that distinguishes between domain, problem-solving, and discourse or communicative actions. By modeling each of these three kinds of actions as separate tree structures, with links between the actions on adjacent levels, our process model enables the incremental recognition of discourse actions that cannot be identified from a single utterance alone. However, it is still able to capture the relationship between discourse, problem-solving, and domain actions. In addition, it provides a more finely differentiated representation of user intentions than previous models, allows the nuances of different kinds of processing (such as different focusing expectations and information retention) to be captured at each level, and accounts for implicit acceptance of a communicated proposition. Our current work involves using this model to handle negotiation dialogues in which a hearer does not automatically accept as valid the proposition communicated by an inform action.

References

- [AP80] James F. Allen and C. Raymond Perrault. Analyzing intention in utterances. *Artificial Intelligence*, 15:143–178, 1980.
- [Car87] Sandra Carberry. Pragmatic modeling: Toward a robust natural language interface. *Computational Intelligence*, 3:117–136, 1987.
- [Car89] Sandra Carberry. A pragmatics-based approach to ellipsis resolution. *Computational Linguistics*, 15(2):75–96, 1989.
- [Gri75] Joseph E. Grimes. *The Thread of Discourse*. Mouton, 1975.
- [GS86] Barbara Grosz and Candace Sidner. Attention, intention, and the structure of discourse. *Computational Linguistics*, 12(3):175–204, 1986.
- [Hov88] Eduard H. Hovy. Planning coherent multisentential text. *Proceedings of the 26th Annual Meeting of the Association for Computational Linguistics*, pages 163–169, 1988.
- [Jos82] Aravind K. Joshi. Mutual beliefs in question-answer systems. In N. Smith, editor, *Mutual Beliefs*, pages 181–197, New York, 1982. Academic Press.
- [LA87] Diane Litman and James Allen. A plan recognition model for subdialogues in conversation. *Cognitive Science*, 11:163–200, 1987.
- [Lam91] Lynn Lambert. Modifying beliefs in a plan-based discourse model. In *Proceedings of the 29th Annual Meeting of the ACL*, Berkeley, CA, June 1991.
- [MP90] Johanna Moore and Cecile Paris. Planning text for advisory dialogues. In *Proceedings of the 27th Annual Meeting of the Association for Computational Linguistics*, pages 203–211, Vancouver, Canada, 1990.
- [MT83] William C. Mann and Sandra A. Thompson. Relational propositions in discourse. Technical Report ISI/RR-83-115, ISI/USC, November 1983.
- [PA80] Raymond Perrault and James Allen. A plan-based analysis of indirect speech acts. *American Journal of Computational Linguistics*, 6(3-4):167–182, 1980.
- [Per90] Raymond Perrault. An application of default logic to speech act theory. In Philip Cohen, Jerry Morgan, and Martha Pollack, editors, *Intentions in Communication*, pages 161–185. MIT Press, Cambridge, Massachusetts, 1990.
- [Pol86a] Livia Polanyi. The linguistics discourse model: Towards a formal theory of discourse structure. Technical Report 6409, Bolt Beranek and Newman Laboratories Inc., Cambridge, Massachusetts, 1986.
- [Pol86b] Martha Pollack. *Inferring Domain Plans in Question-Answering*. PhD thesis, University of Pennsylvania, Philadelphia, Pennsylvania, 1986.
- [Ram89] Lance A. Ramshaw. *Pragmatic Knowledge for Resolving Ill-Formedness*. PhD thesis, University of Delaware, Newark, Delaware, June 1989.
- [Ram91] Lance A. Ramshaw. A three-level model for plan exploration. In *Proceedings of the 29th Annual Meeting of the Association for Computational Linguistics*, Berkeley, California, 1991.
- [Rei78] Rachel Reichman. Conversational coherency. *Cognitive Science*, 2:283–327, 1978.
- [vBC86] Peter van Beek and Robin Cohen. Towards user specific explanations from expert systems. In *Proceedings of the Sixth Canadian Conference on Artificial Intelligence*, pages 194–198, Montreal, Canada, 1986.
- [Wil81] Robert Wilensky. Meta-planning: Representing and using knowledge about planning in problem solving and natural language understanding. *Cognitive Science*, 5:197–233, 1981.