

# The Need for Intentionally-Based Approaches to Language\*

Karen E. Lochbaum  
Aiken Computation Lab  
Harvard University  
Cambridge, MA 02138  
kel@das.harvard.edu

## The Claim

Discourses are inherently intentional; conversational participants engage in them for a reason. Systems for natural language interpretation and generation that do not account, with every utterance, for the purposeful nature of discourse cannot adequately participate in collaborative dialogues.

## The Approach: A SharedPlan Analysis of Subdialogues

Evidence for the above claim comes from my recent work on understanding subdialogues in conversation [Loc93]. Following the work of Grosz and Sidner [GS90], I view discourse behavior as an instance of the more general phenomenon of collaboration. For agents to successfully collaborate on a task, they must hold certain beliefs and intentions regarding the acts they will perform to accomplish that task. The definitions of the SharedPlan model of collaborative activity [GS90, LGS90, GK93] specify the requisite components of the mental states of collaborating agents and consequently provide an important context for interpreting their utterances. In particular, because agents are aware of the beliefs and intentions they must hold to have a SharedPlan, their utterances can be understood as contributing information towards the establishment of the appropriate mental states, and thus the building of such a plan. The process by which utterances are understood has been formalized in algorithms for augmenting the beliefs and intentions of an evolving SharedPlan, and used to explain utterances concerning the performance of actions [LGS90, Loc91].

My current work [Loc93], based on Grosz and Sidner's theory of discourse structure [GS86], provides a new approach to the problem of understanding subdialogues and their relationship to the discourse in which they are embedded. The basic approach entails treating each subdialogue or discourse segment as a separate collaboration between the conversational participants; each utterance of a subdialogue is understood as contributing some information towards the completion of a SharedPlan. Because subdialogues do not occur in isolation, each subdialogue itself is understood in terms of the role its corresponding SharedPlan plays in satisfying the other SharedPlans underlying the dialogue. In particular, if the completion of the new SharedPlan contributes to the establishment of one of the beliefs or intentions necessary for the completion of another, then the first SharedPlan is said to be *subsidiary* to the second. A subsidiary relationship between SharedPlans corresponds to a dominance relationship [GS86] between discourse segment purposes; it provides an explanation [SI81] for why the dominated segment was initiated by a conversational participant<sup>1</sup>.

**An Example** The excerpt in Figure 1, taken from a larger discourse concerned with the replacement of an air compressor pump by an Expert and an Apprentice [GS86], will be used to illustrate the approach. Because the participants in this discourse are collaborating on the replacement of the pump, the first utterance of this excerpt can be understood as establishing mutual belief that the action described therein, i.e. the removal of the flywheel by the Apprentice, is part of the *recipe* the agents will use to accomplish their task [LGS90]. Utterance (2), however, begins a new discourse segment, the purpose of which is to bring it about that the Apprentice knows how to perform (i.e. has a recipe for) the act of removing the flywheel. This segment is recognized as a separate collaboration using a conversational default rule for recognizing an agent's desire to collaborate on

---

\*I would like to thank Barbara Grosz, Christine Nakatani, and Candy Sidner for their comments on this paper. This research has been supported by U S West Advanced Technologies and by a Bellcore Graduate Fellowship.

<sup>1</sup>A detailed description of the model can be found in [Loc93].

the performance of an act [GS90]. The individual utterances of this segment are then understood as directed towards completing a SharedPlan to achieve the purpose of the segment, namely that the Apprentice have a recipe for removing the flywheel. For example, in utterance (3), the Expert's telling the Apprentice the steps in removing the flywheel and an ordering constraint on them (i.e. "First, loosen ..., then pull ...") constitutes a way of achieving that the Apprentice has the requisite recipe.

- (1) E: First you have to remove the flywheel.
- (2) A: *How do I remove the flywheel?*
- (3) E: *First, loosen the two allen head setscrews holding it the shaft, then pull it off.*
- (4) A: *OK.*
- (5) I can only find one screw. Where's the other one?
- (6) E: *On the hub of the flywheel.*

Figure 1: Sample Multi-segment Discourse

To explain the new discourse segment itself, the Expert must determine how the SharedPlan corresponding to the segment is related to the previous SharedPlan. Because the Expert knows that part of having a SharedPlan is mutually believing that the agent of each act is able to perform that act, the Expert can infer that the Apprentice has engaged in the new SharedPlan in order to bring about an instance of that condition required by the previous or "interrupted" SharedPlan. In particular, because utterance (1) has established that removing the flywheel is part of the recipe for replacing the pump, the Expert can infer that the Apprentice has engaged in the subsidiary SharedPlan to bring about a *knowledge precondition* [Mor87] required of the dominating SharedPlan to replace the pump.

## Discussion

Agents engage in subdialogues for many reasons. For example, they may engage in them based on the need for information to perform actions or to weigh options, the need to correct problems that arise during plan execution, or simply as the result of the normal decomposition of a task. My claim is that all subdialogues can be understood in terms of collaborations between agents, and thus modelled using SharedPlans. The type of *plan* does not vary with the type of subdialogue, or an agent's reason for engaging in it, only the object of that plan does<sup>2</sup>. The same methods used in our previous work to understand utterances concerned with "domain goals" [LA87], can also be used to understand utterances aimed at achieving other types of goals; the only difference is in the object of the recipes and SharedPlans used by the algorithms.

**A Rhetorical Relations Approach** This approach contrasts sharply with the less goal-directed account of subdialogues given by Litman and Allen [Lit85, LA87]. In their model, the process of understanding an utterance entails recognizing a *discourse plan* from the utterance and then relating that discourse plan to some *domain plan*. While domain plans represent knowledge about a task, discourse plans represent knowledge about relationships between utterances and plans; for example, an agent may use an utterance to *introduce*, *continue*, or *clarify* a plan. Litman and Allen propose their discourse plans as plan-based correlates of rhetorical relations [Lit86, LA87]. Although these plans address some of the problems with computationally based RST [MT87] analyses (i.e. the formalization and recognition of rhetorical relations), they are extremely rigid in nature and narrow in scope. Each discourse plan (or more specifically the constraints of its decomposition) represents only one specific way in which an utterance can relate to a plan. For example, the only way something can go wrong and be corrected according to the CORRECT-PLAN discourse plan is if a

<sup>2</sup>By the object of the plan, I mean the act on which the agents are collaborating, i.e. the  $\alpha$  in SharedPlan( $\{G_1, G_2\}, \alpha, T_1, T_2$ ).

speaker, not being able to perform the next step in a plan, requests that the hearer do something else first so that he can. In addition, because a new discourse plan is recognized from every utterance, CORRECT-PLAN itself does not model extended problem-solving subdialogues, but only one utterance within such subdialogues. Further utterances of the subdialogue are only understood in terms of their relationship to preceding ones. This approach cannot adequately capture the contribution an utterance of a subdialogue makes to the higher-level purpose of the subdialogue. For example, in the dialogue of Figure 2, utterances (3)-(4) seem intuitively to comprise a subdialogue the purpose of which is to correct a problem that has occurred during plan execution. Utterance (3) identifies the problem, while utterance (4) suggests a way of fixing it. Under Litman and Allen's analysis, however, utterance (3) is understood as an instance of the CORRECT-PLAN discourse plan (with the utterance, the User is correcting the domain plan to add data to a network), while utterance (4) is understood as an instance of IDENTIFY-PARAMETER. The parameter that utterance (4) is understood to be identifying is one in the CORRECT-PLAN discourse plan, namely the parameter that specifies what new step is being added to a domain plan to correct it. Not only does this analysis run counter to intuitions as to what utterances (3) and (4) both individually and collectively are about, but, as used in a model of plan recognition, it constitutes a claim that the speaker (i.e. User) (i) produces utterances (3) and (4) intending to perform the above CORRECT-PLAN and IDENTIFY-PARAMETER actions respectively, and (ii) intends that the hearer (System) recognize these intentions.

- (1) User: Show me the generic concept called "employee".
- (2) System: OK. <system displays network>
- (3) User: *I can't fit a new ic below it.*
- (4)       *Can you move it up?*
- (5) System: Yes. <system displays network>
- (6) User: OK, now make...

Figure 2: A Correction Subdialogue (taken from Litman[Lit85])

The problem with Litman and Allen's approach, like RST-based approaches in general, is that it essentially provides only an utterance-to-utterance based analysis of discourse. In addition to not recognizing discourse segments as separate units with an overall purpose, the model also fails to recognize a subdialogue's relationship to the discourse in which it is embedded. That is, it cannot account for *why* agents engage in subdialogues. More recent models [LC91, LC92, Ram91] that augment Litman and Allen's two types of plans with other types also suffer from the same shortcomings<sup>3</sup>.

**Evidence from Generation** Work in generation has recognized a similar problem with respect to RST-based approaches. In particular, Moore & Paris [MP91] (see also [MP92, Hov93]) have argued for the need to augment RST-based text plans or schemas [Hov88, McK85] with an intentional structure in order to respond to follow-up questions. The problem is that although solely RST-based approaches associate a communicative goal with each schema, they do not represent the intended effect of each component of the schema, nor the role that each component plays in satisfying the overall communicative goal associated with the schema. Without such information, a system cannot respond effectively if the hearer does not understand or accept its utterances. In response to this problem, Moore and Paris have devised a planner that constructs text plans containing both intentional and rhetorical information. By recording these text plans as part of the dialogue history, their system is able to reason about its previous utterances in interpreting and responding to users' follow-up questions.

**Conclusions** Both the interpretation process and the generation process need intentionally-based approaches to language. In the former, a solely intentional approach provides a more general

<sup>3</sup>Analyses of all of these approaches can also be found in [Loc93].

model for understanding subdialogues and their relationships. In the latter, intentional information augments RST information to allow more effective participation in explanation dialogues. Although rhetorical relations have proved useful in machine-based natural language generation (see Hovy's recent survey [Hov93]), their cognitive role remains unclear. Does a speaker actually have them "in mind" when he produces utterances? Or are they only "compilations" of intentional information that are computationally efficient for generation systems [MP91]? And if a speaker does have rhetorical relations in mind, does a hearer actually infer them? On that matter, I'd argue, based on the above discussion (and following Grosz and Sidner [GS86]), that a discourse can be understood even if the hearer (be it machine or person) cannot infer, construct, or name any such relations used by the speaker.

## References

- [GK93] B. J. Grosz and S. Kraus. Collaborative plans for group activities. In *Proceedings of IJCAI-93*, Chambéry, Savoie, France, 1993.
- [GS86] B.J. Grosz and C.L. Sidner. Attention, intentions, and the structure of discourse. *Computational Linguistics*, 12(3), 1986.
- [GS90] B.J. Grosz and C.L. Sidner. Plans for discourse. In P.R. Cohen, J.L. Morgan, and M.E. Pollack, editors, *Intentions in Communication*. MIT Press, 1990.
- [Hov88] E. H. Hovy. Planning coherent multisentential text. In *Proceedings of ACL-88*.
- [Hov93] E.H. Hovy. Automated discourse generation using discourse structure relations. *Artificial Intelligence*, 1993. To appear.
- [LA87] D.J. Litman and J.F. Allen. A plan recognition model for subdialogues in conversations. *Cognitive Science*, 11:163-200, 1987.
- [LC91] L. Lambert and S. Carberry. A tripartite plan-based model of dialogue. In *Proceedings of ACL-91*.
- [LC92] L. Lambert and S. Carberry. Modeling negotiation subdialogues. In *Proceedings of ACL-92*.
- [LGS90] K. E. Lochbaum, B. J. Grosz, and C. L. Sidner. Models of plans to support communication: An initial report. In *Proceedings of AAAI-90*, Boston, MA, 1990.
- [Lit85] D. J. Litman. *Plan Recognition and Discourse Analysis: An Integrated Approach for Understanding Dialogues*. PhD thesis, University of Rochester, 1985.
- [Lit86] D. J. Litman. Linguistic coherence: A plan-based alternative. In *Proceedings of ACL-86*.
- [Loc91] K. E. Lochbaum. An algorithm for plan recognition in collaborative discourse. In *Proceedings of ACL-91*.
- [Loc93] K. E. Lochbaum. A collaborative planning approach to understanding subdialogues in conversation. Technical report, Harvard University, 1993.
- [McK85] K. R. McKeown. *Text Generation: Using Discourse Strategies and Focus Constraints to Generate Natural Language Text*. Cambridge University Press, Cambridge, England, 1985.
- [Mor87] L. Morgenstern. Knowledge preconditions for actions and plans. In *Proceedings of IJCAI-87*.
- [MP91] J. D. Moore and C. L. Paris. Discourse structure for explanatory dialogues. In *AAAI Fall 1991 Symposium on Discourse Structure in Natural Language Understanding and Generation*, Asilomar.
- [MP92] J.D. Moore and M.E. Pollack. A problem for RST: The need for multi-level discourse analysis. *Computational Linguistics*, 18(4), December 1992.
- [MT87] W.C. Mann and S.A. Thompson. Rhetorical structure theory: A theory of text organization. In L. Polanyi, editor, *The Structure of Discourse*. Ablex Publishing Corp., 1987.
- [Ram91] L. A. Ramshaw. A three-level model for plan exploration. In *Proceedings of ACL-91*.
- [SI81] C.L. Sidner and D. J. Israel. Recognizing intended meaning and speakers' plans. In *Proceedings of IJCAI-81*.