

A PRAGMATICS-BASED APPROACH TO ELLIPSIS RESOLUTION

Sandra Carberry

Department of Computer and Information Sciences
University of Delaware
Newark, DE 19716

Intersentential elliptical utterances occur frequently during information-seeking dialogs in task domains. This paper presents a pragmatics-based framework for interpreting such utterances. Discourse expectations and focusing heuristics are used to facilitate recognition of an information-seeker's intent in uttering an elliptical fragment. The ellipsis is comprehended by identifying both the aspect of the information-seeker's task-related plan highlighted by the fragment and the conversational discourse goal fulfilled by the utterance. The contribution of this approach is its consideration of pragmatic information, including discourse content and conversational goals, rather than just the precise representation of the preceding utterance.

1 INTRODUCTION

Incomplete utterances are common in communication between humans. They range from sentences that fail to include all requisite semantic information to syntactically incomplete sentence fragments. In many cases, these utterances cannot be understood in isolation, but must be interpreted within the established context. Precisely how this should be done is a difficult problem for natural language systems.

One might suggest that the problem be avoided in man-machine communication by training human users to employ only syntactically and semantically complete utterances. However, this does not appear to be feasible, as demonstrated in an empirical study conducted by Carbonell (1983) in which it was shown that human users find it easy to avoid complex syntactic structures but difficult to avoid incomplete utterances.

Even if it were possible to train users to avoid incomplete utterances, these restrictions would be undesirable. Constraining man-machine communication to only a subset of the utterances normally employed by humans would force users to give less attention to their problem solving goals in order to concentrate more on the preciseness of their utterances. In addition, it appears that fragmentary utterances are not merely a result of sloppy communication. Although every utterance has a **discourse goal** (a conversational or communicative goal such as answering a question or seeking clarification), elliptical fragments are often used to accomplish discourse goals that would require more

effort to convey with a complete sentence. For example, in the following dialog¹, Speaker 2's fragment expresses doubt about the proposition stated by Speaker 1.

Example 1

Speaker 1: "The Korean jet shot down by the Soviets was a spy plane."

Speaker 2: "With 269 people on board?"

However, a complete sentence such as

Speaker 2: "Was the Korean jet shot down by the Soviets a spy plane with 269 people on board?"

fails to adequately communicate the doubt conveyed by the previous fragment. Only a more complex sentence that marks the discourse goal, such as

Speaker 2: "How can you think that the Korean jet shot down by the Soviets was a spy plane, when it had 269 people on board?"

will accomplish this objective.

Thus a robust natural language interface must handle the kinds of incomplete utterances normally used by humans. To do otherwise is to prohibit communication that humans regard as natural, and therefore detract from their ability to communicate as effectively with machines as they do with one another.

Contextual ellipsis in dialog is the use of a sentence fragment (a syntactically incomplete utterance), along with the context established by the preceding dialog, to communicate a complete thought and accomplish a

Copyright 1989 by the Association for Computational Linguistics. Permission to copy without fee all or part of this material is granted provided that the copies are not made for direct commercial advantage and the *CL* reference and this copyright notice are included on the first page. To copy otherwise, or to republish, requires a fee and/or specific permission.

0362-613X/89/010075-96-\$03.00

speech act². Such fragments are often referred to as **intersentential ellipsis** since they appear between sentences in a dialog. It can be seen, however, that understanding intersentential ellipsis often depends more on pragmatic knowledge, such as the inferred task-related plan and discourse goals motivating the speaker, than on the syntactic structure or semantic content of preceding utterances. This is illustrated by Examples 2 and 3, in which Speaker 1's communicated goal and the relevant plans for accomplishing it play a major role in understanding the elliptical fragment.

Example 2

Speaker 1: "I want to cash this check. Small bills only, please."

Example 3³

Speaker 1: "Who are the candidates for programming consultants?"

Speaker 2: "Mary Smith, Bob Jones, and Ann Doe have applied for the job."

Speaker 1: "Tom's recommendation?"

Speaker 2: "He thinks Bob Jones and Ann Doe have the necessary background and should be invited for an interview."

In Example 2, Speaker 1's goal is to cash a check, and relevant plans can include a subplan for getting a particular distribution of paper money; in Example 3, Speaker 1's goal is hiring programming consultants, perhaps for an introductory programming course, and the relevant plans for doing this include a subplan for identifying the best applicant.

Previous research on understanding intersentential ellipsis has emphasized syntactic and semantic strategies (Hendrix et al. 1978; Waltz 1978; Weischedel and Sondheimer 1982; Carbonell 1983), but the contributions of the speaker's plans and goals to the interpretation of ellipsis has hitherto been inadequately explored. One objective of our research has been to investigate a plan-based framework for understanding intersentential ellipsis that occurs in task-oriented dialogs. Our work shows that elliptical fragments can be viewed as highlighting aspects of the information-seeker's task-related plan, with the focus of attention in the plan providing the context in which the fragment should be understood.

But identifying the aspect of the speaker's plan highlighted by the ellipsis is not enough to understand the utterance. As shown by Grosz (1979), a speaker may pursue several different kinds of goals with a single utterance. For example, an information-seeker has the long-term goal of performing a task comprised of sub-tasks; this hierarchical structure produces a set of task-related goals. The information seeker is attempting to construct a plan for his underlying task, resulting in plan formation goals. In addition, each utterance pursues a more immediate discourse goal, such as requesting information or seeking clarification. As has been shown by several researchers (Mann et al. 1977; Reich-

man 1978; Pollack et al. 1982), one must determine not only the element of the task being addressed by an utterance, but also the discourse goal being fulfilled by it. Previous research on understanding ellipsis has ignored this aspect. Thus a second objective of our research has been to formulate a strategy for recognizing discourse goals communicated by elliptical fragments. Our work indicates that mutual beliefs and expectations developed from the preceding dialog play a major role in identifying the intent behind intersentential ellipsis.

Our processing framework coordinates many knowledge sources, including discourse expectations, inferred beliefs, the information-seeker's inferred task-related plan, and focusing strategies, to produce a rich interpretation of ellipsis. It is the first interpretation strategy to address the problem of identifying the discourse goal accomplished via an intersentential elliptical fragment. As a result, this pragmatics-based framework understands elliptical fragments that other systems cannot handle.

2 STRATEGIES FOR INTERPRETING INTERSENTENTIAL ELLIPSIS

As illustrated by Examples 1, 2, and 3, intersentential elliptical fragments cannot be understood in isolation from the context in which they occur. Therefore a strategy for interpreting such fragments must rely on knowledge obtained from sources other than the fragment itself. Three possibilities exist: the syntactic form of preceding utterances, the semantic representation of preceding utterances, and expectations gleaned from understanding the preceding discourse.

A strategy using the first kind of knowledge might perform two transformations on the syntactic form of a preceding utterance.

1. a substitution/expansion operation in which the elliptical fragment is substituted for a syntactic constituent in the preceding utterance, or the syntactic representation of the preceding utterance is expanded to accommodate the fragment.
2. a transformation operation that maps questions into answers, statements into questions, etc.

Syntactic strategies are exemplified by the work in Hendrix et al. (1978) and Weischedel et al. (1982). However, many examples of ellipsis cannot be handled by syntactic strategies. Consider again the dialog presented in Example 2. No transformation or modification of Speaker 1's first statement will produce an utterance that represents the intended meaning of Speaker 1's elliptical fragment.

The second potential strategy uses the semantic representation of the last utterance as a pattern that suggests slots for which an elliptical fragment may provide a filler or a substitution. Semantic strategies are presented in Waltz (1978) and Carbonell (1983). However, semantic strategies require an extensive case

frame representation and are inadequate for handling elliptical fragments that rely on an assumed communication of the underlying task. For example, it is improbable that the case frame for programming consultant in Example 3 would contain slots for the recommendations of individual faculty.

The third potential strategy utilizes pragmatic knowledge, such as a model of the information-seeker's inferred task-related plan and expected discourse goals. The power of this approach is its reliance on knowledge gleaned from the dialog, including discourse content and conversational goals, rather than on precise representations of the preceding utterances alone.

Allen (Allen et al. 1980) addressed the problem of understanding sentence fragments occurring at the outset of a dialog, and was the first to relate ellipsis processing to the domain-dependent plan underlying a speaker's utterance. Allen viewed the speaker's utterance as part of a plan that the speaker had constructed and was executing to accomplish his overall task-related goals (in this case, to meet or board a train). To interpret elliptical fragments, Allen first constructed a set of possible surface speech act representations for the elliptical fragment, limited by syntactic clues appearing in the fragment; he then attempted to infer the speaker's goal-related plan that included an execution of the observed utterance. A part of this inference process involved determining which of the partially constructed plans, connecting the task-related goals of meeting or boarding a train and the observed utterance, were reasonable given the knowledge and mutual beliefs of the speaker and hearer. Allen selected the surface speech act that produced the most reasonable inferred plan as the correct interpretation.

Allen noted that the speaker's fragment must identify the subgoals that the speaker is pursuing, but claimed that in very restricted domains, identifying the speaker's overall goal from the utterance was sufficient to identify the appropriate response in terms of the obstacles present in such a plan. For isolated sentence fragments in his restricted domain of train arrivals and departures, Allen's interpretation strategy worked well. To understand elliptical fragments during a dialog in a more complex domain, it is necessary to identify the particular aspect of the speaker's overall task-related plan addressed by the fragment and recognize the discourse or communicative goal being pursued.

More recently, Litman (1986) developed a plan-based theory that incorporates both domain plans and metaplans to understand utterances. Similar to Allen, Litman used syntactic clues in the utterance to postulate speech acts "request" or "inform" that an elliptical fragment was intended to accomplish, and she then attempted to construct an inference path from a postulated speech act to the plan inferred for the speaker. However, it appears that Litman's system would handle an elliptical fragment such as

"With 269 people on board?"

from Example 1 in much the same way as the complete sentence

"Did the Korean jet have 269 people on board?"

As such, Litman's strategy would be unable to recognize the surprise and doubt conveyed by the elliptical fragment in Example 1. In addition, Litman's metaplans (such as **Introduce-Plan**, **Continue-Plan**, and **Modify-Plan**) are more like plan construction metaplans than metaplans representing communicative goals such as expressing surprise or seeking clarification of a question. We contend that recognition of the speaker's discourse or communicative goal must be an integral part of any strategy that really understands ellipsis.

In addition to the syntactic, semantic, and plan-based strategies, a few other heuristics have been devised. Carbonell (1983) used rules that suggested a set of expected user utterances and related elliptical fragments to these expected patterns. For example, if the system asked the user whether a particular value should be used as the filler of a slot in a case frame, the system then expected the user's utterance to contain a confirmation or disconfirmation pattern, a different filler for the slot, a comparative pattern such as "too hard", and so forth. Although these rules use expectations about how the speaker might respond, they seem to have little to do with the expected discourse goals of the speaker; as a result, they would be unable to differentiate fragments that merely request information from fragments whose intent is to seek clarification or express surprise.

3 SCOPE OF THE RESEARCH

This paper describes a pragmatics-based approach for understanding intersentential elliptical fragments that occur in information-seeking dialogs. We have studied dialogs in which the information-seeker is attempting to construct a plan for accomplishing a task, but the plan itself will be executed at some time in the future. Examples of such tasks include expanding a company's product line, purchasing a home, pursuing a degree, or even taking a vacation. We have selected this class of dialogs because they are typical of a large percentage of the interactions with database management systems, decision support systems, and expert systems. However, the principles presented in this paper should be extendible to other kinds of dialogs.

Our hypotheses and strategies were motivated by an analysis of naturally occurring dialogs from several domains, including a radio call-in show⁴ providing advice on investments, interactions with the REL (Thompson 1980) natural language interface to a ship data base, and student advisement sessions. We are assuming that the natural language system plays the role of information-provider and that communication is via a typical terminal. Since our primary interest in this research is the affect of expectations and inferred knowledge on

ellipsis understanding, we have not considered the contribution of clue words (words and phrases, such as "Now" and "As I was saying", that contain information about the structure of a discourse (Reichman 1978)) to the processing of ellipsis. However, they certainly could be incorporated into our strategy.

In this research, we are attempting to ascertain the extent to which pragmatic knowledge can be used to understand ellipsis and the limitations of such an approach. However, we do not maintain that a natural language system should use only pragmatic knowledge; a robust system will need to coordinate syntactic, semantic, and pragmatic techniques in order to fully understand the wide variety of elliptical utterances employed in human communication.

In the remainder of this paper, we argue for the use of pragmatic knowledge in ellipsis understanding, and present a pragmatics-based approach for understanding intersentential elliptical fragments uttered by a person seeking information in order to construct a plan for accomplishing a task. Section 4 discusses the knowledge needed to understand elliptical fragments. Section 5 describes our overall interpretation framework, which is presented in detail in sections 6–8. Our strategy views elliptical fragments as highlighting, or bringing to attention, aspects of the information-seeker's underlying task-related plan, with discourse expectations used to guide interpretation. Section 6 presents the discourse or communicative goals that we have found pursued via elliptical fragments in information-seeking dialogs. Section 7 discusses how discourse expectations develop during a dialog, describes how these expectations suggest discourse goals that an information-seeker might pursue, and presents discourse goal rules for recognizing elliptical fragments that accomplish a particular discourse goal. The discourse goal rules guide how the fragment's highlighting of a component of the task-related plan should be understood; Section 8 describes a strategy for identifying the highlighted portion of the plan. Section 9 illustrates our framework with several examples, and sections 10 and 11 discuss future extensions to this work and our conclusions. Throughout this paper, the information seeker and information provider will be referred to as **IS** and **IP** respectively.

4 REQUISITE KNOWLEDGE

A speaker can felicitously employ intersentential ellipsis only if he believes his utterance will be properly understood and his intent recognized by the listener. The motivation for this work is the hypothesis that speaker and hearer mutually believe that certain knowledge has been acquired during the course of the dialog and that this factual knowledge along with other processing knowledge will be used to deduce the speaker's intentions. We claim that the requisite factual knowledge includes the speaker's inferred task-related plan, the speaker's inferred beliefs, and the anticipated dis-

course goals of the speaker; we claim that the requisite processing knowledge includes plan recognition strategies and focusing techniques.

4.1 TASK-RELATED PLAN

According to Grice (1957,1969), a listener must infer a speaker's intent in making an utterance. Now consider an incomplete utterance. The listener, in attempting to deduce the speaker's intent, will be guided by Grice's conversational maxims (Grice 1975). In particular, Grice's maxims of manner and relation suggest that the speaker believes his utterance is an adequate vehicle for conveying his intentions and that the utterance is relevant to the current dialog. A cooperative listener will have assimilated the preceding dialog, inferred much of the underlying task-related plan motivating the speaker's queries, and be focused on that aspect of the task on which the information-seeker's attention was centered (Grosz 1977; Carberry 1983; Carberry 1988) immediately prior to the new utterance. Given an incomplete utterance, the listener can use this acquired knowledge to attempt to deduce the speaker's intentions by trying to fit the utterance into the partially constructed plan, and thereby enable the dialog to continue without interruption.

This claim that the speaker's underlying task-related plan is an essential component of an ellipsis interpretation strategy is further supported by research demonstrating the need for considering such plans in understanding other types of utterances. Grosz (1977) and Robinson (1981) used models of the speaker's task in apprentice-expert dialogs to determine the referents of definite noun phrases and verb phrases. Perrault and Allen (Perrault et al. 1980) used expectations of speaker goals, and inference of their plans for achieving these goals, in the interpretation of indirect speech acts. In addition, Allen (Allen et al. 1980) introduced the concept of a plan-based strategy for interpreting fragmentary utterances at the outset of a dialog in a restricted domain.

4.2 SHARED BELIEFS

Shared beliefs, beliefs which the listener believes speaker and listener mutually hold, are a second component of factual knowledge required for processing intersentential elliptical fragments. These shared beliefs either represent presumed a priori knowledge of the domain, such as a presumption that dialog participants in a university domain know that each course has a teacher, or beliefs derived from the dialog itself. An example of the latter occurs if IP tells IS that CS360 is a 5 credit hour course; IS may not himself believe that CS360 is a 5 credit hour course, but as a result of IP's utterance and the assumption of a cooperative dialog, he does believe it is mutually believed that IP believes this.

Real understanding requires that we identify the speaker's discourse goal in making the utterance.

Shared beliefs, often called mutual beliefs, form a part of communicated knowledge used to interpret utterances and identify discourse goals in a cooperative dialog. The following two dialog sequences illustrate how IP's beliefs about IS influence understanding.

IS: "When does CS400 meet?"
 IP: "CS400 meets on Monday from 7 pm until 9 pm."
 IS: "Who is teaching it?"
 IP: "Dr. Brown."
 IS: "At night?"

Most people would interpret the fragmentary utterance "At night?" as an expression of surprise and request for corroboration or explanation. Now consider the following dialog:

IS: "Who is teaching CS400?"
 IP: "Dr. Brown is teaching CS400."
 IS: "At night?"

In this case, the elliptical fragment "At night?" would be understood as a simple request to know whether CS400 is meeting at night. The reason for this difference in interpretation is the difference in beliefs regarding IS at the time the elliptical fragment is uttered. In the latter case, IP believes it is mutually believed that IS does not know IP's beliefs about the meeting time of CS400. Since IS cannot be surprised about something he doesn't know, a different intention or discourse goal is attributed to him.

Allen and Perrault (Perrault et al. 1980) used mutual beliefs in their work on indirect speech acts and suggested their use in clarification and correction dialogs. Sidner (1983, 1985) models user beliefs about system capabilities in her work on recognizing speaker intention in utterances. Mutual beliefs appear to be a major influence on how human listeners interpret and respond to utterances.

4.3 DISCOURSE EXPECTATIONS

A number of researchers (Sacks et al. 1974; Labov and Fanshel 1977; Mann et al. 1977) have noted regularities in dialog and have attempted to capture these regularities in knowledge structures that can be used for comprehending utterances. For example, Mann, Moore, and Levin (Mann et al. 1977) used a knowledge structure called a *dialog game* to model goal-related use of language in joint interactions such as buying/selling and learning/teaching. Grosz and Sidner (Grosz et al. 1986) and Reichman (1984) have investigated discourse structure and have shown that a coherent discourse can be segmented into units that have well-defined relationships to one another. Reichman contended further that the existing discourse structure established expectations about appropriate next conversational moves.

Our analysis of naturally occurring dialog indicates that such expectations about appropriate next steps in the dialog form a third component of factual knowledge that plays a major role in comprehending elliptical

fragments. The dialog preceding an elliptical utterance may suggest discourse or communicative goals for the speaker, such as seeking clarification of a question, expressing surprise at a question, or answering a question. Our transcript analysis leads us to hypothesize that these suggested discourse goals become shared knowledge between speaker and hearer and that, as a result, the listener is on the lookout for the speaker to pursue these anticipated discourse goals and interprets utterances accordingly.

Consider for example the following dialog:

Example 4

IP: "Do you want to take CS360?"
 IS: "Who is teaching it?"
 IP: "Dr. Brown and Dr. White."
 IS: "Yes, with Dr. White."

In this example, IP's initial query produces a strong anticipation that IS will pursue the discourse goal of providing the requested information. Therefore subsequent utterances should be processed with the expectation that IS will eventually address this goal. IS's first utterance is interpreted as pursuing a discourse goal of seeking additional information in order to answer the question. IS's last utterance reverts back to addressing the original expectation and answers the initial query posed by IP.

4.4 PROCESSING KNOWLEDGE

Our transcript analysis indicates that plan recognition strategies and focusing techniques are necessary components of processing knowledge for interpreting intersentential ellipsis. Plan recognition strategies are essential in order to infer a model of the speaker's underlying task-related plan, shown earlier to be an essential component of factual knowledge, and to consider possible expansions of that plan when processing elliptical fragments.

Focusing techniques (Grosz 1977; Sidner 1981; Robinson 1981; McKeown 1985; Carberry 1988) are necessary in order to identify that portion of the underlying plan to which a fragmentary utterance refers. Consider again the dialog in Example 3. The focus of attention in this dialog is on considering the job applicants and evaluating them; therefore IS is most expected to continue with this subtask before considering other aspects of his overall task (Carberry 1983, 1988). As a result, IS's fragmentary utterance "Tom's recommendation?" will be understood as a request for Tom's opinion about the suitability of the job applicants for the job of programming consultant. However, if the focus of attention were instead on determining the number of consultants needed, then IS's fragmentary utterance would be understood differently.

5 FRAMEWORK FOR PROCESSING ELLIPSIS

If an utterance is parsed as a sentence fragment, ellipsis processing begins. In our approach, we assume that the

processor has access to a model of the preceding dialog containing 1. a context tree (Carberry 1988) representing IS's inferred underlying task-related plan; 2. a stack containing the discourse expectations for IS, and 3. a belief model representing IS's inferred beliefs. These comprise the requisite factual knowledge discussed in the preceding section.

We claim that the discourse expectations for IS should be used to guide interpretation of fragmentary utterances. The top element of the discourse stack represents the most immediate discourse expectation for IS and suggests potential discourse goals that IS might be expected to pursue. If IS employs a complete utterance, the utterance generally gives clues about IS's discourse goal; in this case, the discourse expectations play a secondary role. However, if IS employs an elliptical fragment, then he must intend that the fragment be interpreted according to these expectations. IS believes that IP expects him to pursue certain discourse goals and therefore that IP will be anticipating utterances directed toward these goals. If IS does not want his utterance interpreted according to these expectations, then his utterance must override them, either by not producing an interpretation appropriate to an expected discourse goal or by explicitly suggesting a different discourse goal (via the use of clue words). Otherwise IS is assumed to *intend* that his utterance be interpreted according to IP's expectations.

The following alternative utterances by IS illustrate this hypothesis:

IP: "Do you want to take CS360?"

Response 1: IS: "Are you asking me if I want to take it next semester?"

Response 2: IS: "Will it be offered next semester?"

Response 3: IS: "Next semester?"

In Response 1, IS's utterance indicates that it is an attempt to clarify the question posed by IP; in Response 2, IS's utterance simply requests information about CS360 in order to formulate an answer to IP's question. However, in Response 3, IS's utterance is an elliptical fragment that conceivably could produce several different interpretations. It might be an attempt to clarify the question posed by IP, it might be a request for information about when CS360 is offered in order to formulate an answer to the posed question, or it might be an expression of surprise that IP would ask such a question (for example, if CS360 is only offered next semester and IS already has a very full schedule). Such elliptical fragments do not explicitly indicate IS's discourse goal in uttering the fragment and thus understanding relies heavily on mutual beliefs, including the discourse goals that IS is expected to pursue.

Thus we claim that IS must intend that his elliptical fragment be interpreted according to these mutually believed expectations. Figure 1 presents our algorithm for ellipsis resolution. It uses the stack of discourse expectations to control interpretation of elliptical frag-

ments (Loop A of Figure 1). Associated with each expectation is a discourse expectation rule that suggests discourse goals that IS might pursue (steps 2 and 3 in Figure 1). Each suggested discourse goal has one or more associated discourse goal rules that are applied in an attempt to produce a coherent interpretation of the elliptical fragment (Loop B in Figure 1). We view IS's elliptical fragment as highlighting, or bringing to attention, some aspect of the underlying task-related plan. The condition part of a discourse goal rule uses the aspect of the plan highlighted by the elliptical fragment, as identified by a plan analyzer (step 9 in Figure 1), along with other factual and processing knowledge to determine whether its associated discourse goal should be recognized as the discourse goal being pursued by IS. If the condition part of a discourse goal rule is satisfied, the rule then produces an interpretation of the elliptical fragment relevant to the recognized discourse goal (steps 11–13 in Figure 1). The individual steps of the algorithm in Figure 1 will be described in the following sections.

6 DISCOURSE GOALS

We define a discourse goal intuitively as what the speaker is trying to do in making an utterance. The discourse goal itself is content independent, although its realization in an utterance will involve some term or proposition from the speaker's underlying task-related plan. For example, the discourse goal Express-Surprise-Obtain-Corroboation represents a speaker's attempt to convey to his listener that he is surprised at some proposition P and would like justification of it.

A number of researchers have contended that a coherent discourse consists of segments that are related to one another. The structuring relations have been given many names, including rhetorical predicates (Grimes 1975, McKeown 1985), coherence relations (Hobbs 1979), discourse segment purposes (Grosz 1986), and conversational moves (Reichman 1984). Our concept of discourse goals is closest to Grosz and Sidner's discourse segment purposes and Reichman's conversational moves. Grosz and Sidner define a discourse segment purpose as the intended role of a discourse segment in achieving an overall discourse purpose. They contend that the discourse segment successfully serves its intended function only if the discourse segment purpose is recognized by the listener. Similarly, we contend that understanding an elliptical fragment requires that the listener identify the discourse goal being pursued by the speaker in uttering the fragment. Reichman differentiated between utterances that continued the current discourse segment and utterances that constituted conversational moves to a segment playing a different role in the overall dialog. Some of her conversational moves, such as challenging a claim, had the flavor of goals that one might pursue in a dialog. Similar to our work, Reichman contended that

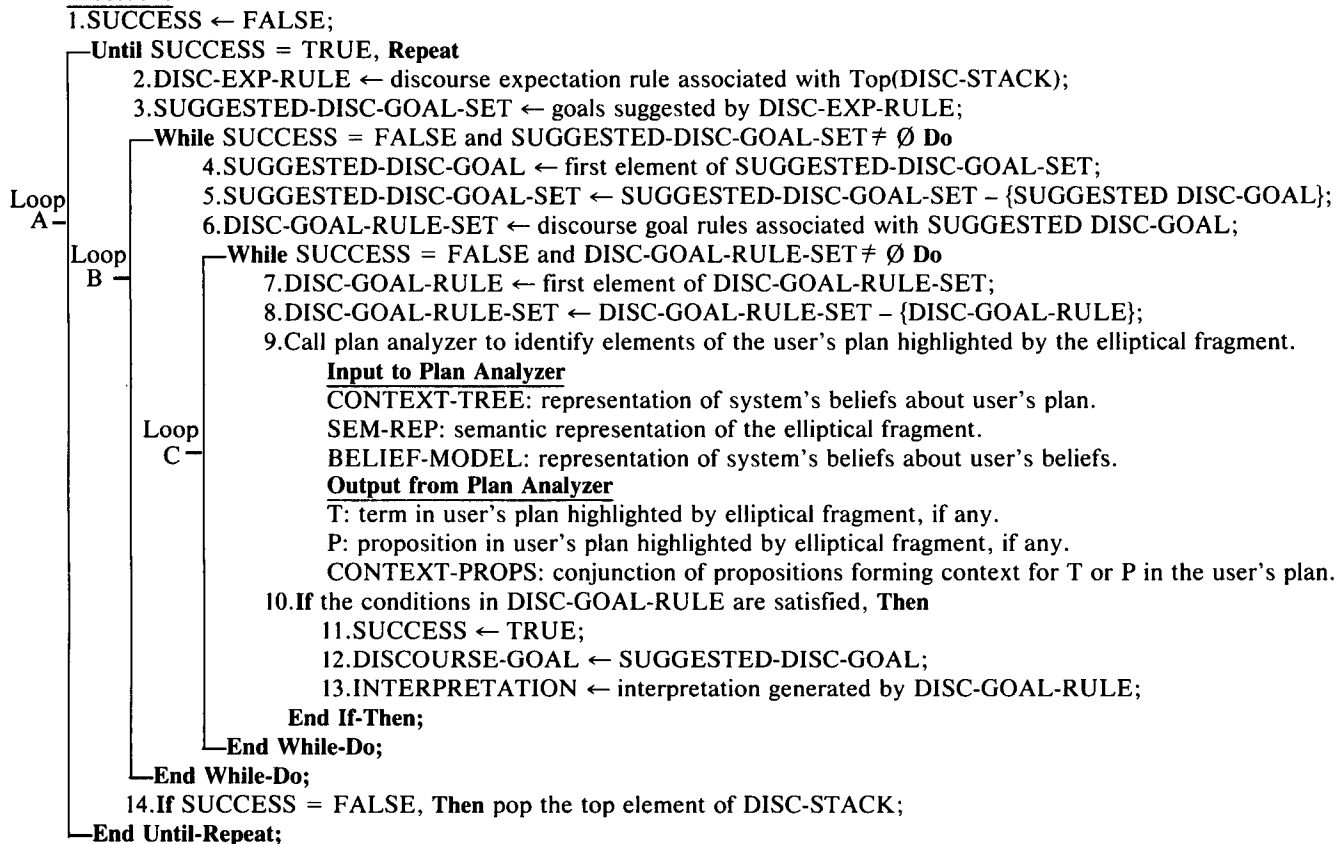
Input

DISC-STACK: discourse stack containing expectations about the user's discourse behavior, along with the semantic representation of the utterance that prompted each expectation.

SEM-REP: semantic representation of the elliptical fragment.

CONTEXT-TREE: representation of the system's beliefs about the user's underlying task-related plan.

BELIEF-MODEL: representation of the system's beliefs about the user's beliefs.

Procedure**Output**

DISCOURSE GOAL: user's identified discourse goal

INTERPRETATION: Interpretation of the elliptical fragment

Figure 1. Algorithm for Interpreting Elliptical Fragments.

conversational moves establish expectations about subsequent conversational moves. However, both discourse segment purposes and conversational moves differ from our discourse goals in that a discourse segment may contain more than one utterance—ie., discourse segment purposes and conversational moves relate discourse segments, not utterances. In addition, neither Grosz and Sidner nor Reichman provide the details of a computational mechanism for identifying the role of an utterance in a dialog. Our theory does this for elliptical fragments in an information-seeking dialog. Given the discourse goal expectations, it provides a mechanism for understanding an elliptical fragment by recognizing the intentions it communicates.

We have analyzed dialogs from several different

domains and have identified 15 discourse goals that occur during information seeking dialogs and which may be accomplished via elliptical fragments. The following are some of these discourse goals, with illustrative examples.

1. **Provide-For-Assimilation**: IS provides information pertinent to formulation of his underlying task-related plan.

IS: "I want to get a degree."
"CS major."

2. **Obtain-Information**: IS requests information relevant to constructing the underlying task-related plan or relevant to formulating an answer to a question posed by IP.

IS: "Is CS360 being offered this fall?"

IP: "Yes."

IS: "The instructor?"

3. **Express-Surprise-Obtain-Corroboration:** IS expresses surprise regarding some proposition P and requests elaboration on, and justification of, it.

IS: "What time does CS360 meet this fall?"

IP: "Monday, Wednesday, and Friday at 8 am."

IS: "Who is teaching it?"

IP: "Dr. Smith."

IS: "At 8 am?"

4. **Seek-Identify:** IS is unable to satisfactorily identify the referent of an item in IP's utterance and requests help from IP in doing so.

IS: "What is Dr. Smith teaching this fall?"

IP: "CS360."

IS: "The course in architecture?"

5. **Seek-Clarify-Question:** IS requests information relevant to clarifying a question posed by IP.

IP: "Do you want to take CS105?"

IS: "Next semester?"

6. **Suggest-Answer-Own-Question:** IS suggests an answer to his own question. IS's intent is to suggest that IP give particular consideration to IS's proposal in formulating an answer to the question.

IS: "What course should I take during winter session?"

"CS370?"

7. **Answer-Question-With-Restrictions:** IS answers a Yes-No question, providing restrictions on the relevant underlying task-related plan.

IP: "Would you like to take CS360?"

IS: "At night with Dr. White."

Although they are not explicitly accomplished via elliptical fragments, two other discourse goals play a major role in understanding ellipsis and will be discussed further in the next section:

- **Accept-Response:** IP has responded to IS's request for information; IS accepts the response.
- **Accept-Question:** IP has posed a question to IS; IS accepts the question.

7 DISCOURSE COMPONENT

7.1 DISCOURSE EXPECTATIONS

The discourse goals of the preceding section also serve as discourse expectations. As argued in Section 4.3, discourse expectations play a major role in the comprehension of elliptical fragments. When IS makes an utterance, he is attempting to accomplish a discourse goal; this discourse goal may in turn establish expectations about what IS will do next. For example, if IS asks a question, one anticipates that IS may want to expand on his question by further identifying an entity or by clarifying the question.

Similarly, utterances made by IP also establish expectations for IS. These discourse expectations may

often be met implicitly as well as explicitly. Consider what happens when IP poses a question to IS. In a cooperative dialog, we expect IS to answer the question. But before that can occur, IS must understand the question and accept it as relevant and valid. Normally dialog participants accept such questions implicitly by answering the question or by proceeding to seek information relevant to formulating an answer. However, IS may be unable to accept the question posed by IP because he does not understand it (perhaps he is unable to identify some of the entities mentioned in the question), or because he is surprised by it. This leads IS to pursue discourse goals such as seeking confirmation, seeking the identify of an entity, seeking clarification of the posed question, or expressing surprise at the question. Thus, when IP poses a question to IS, our discourse expectations are that IS will first accept the question (or if that is not yet possible, work towards accepting it) and then answer it. Thus these two expectations must be pushed onto the discourse stack and should serve to guide our understanding of elliptical fragments.

7.2 THE DISCOURSE STACK

The discourse stack contains expectations about IS's discourse behavior, along with the semantic representation of the utterance that prompted the expectation. We contend that understanding elliptical fragments relies on discourse expectations that are pushed onto or popped from the stack as a result of utterances made by IS and IP. Precisely how this is accomplished in all cases is unclear. However, in this section, we will motivate the contents of the discourse stack by giving a set of stack processing rules that hold for simple utterances⁵. These rules presume that the discourse goal accomplished by an utterance can be identified. Although our theory for ellipsis processing includes recognition of the discourse goal fulfilled by an elliptical fragment, we have not investigated the recognition of discourse goals pursued via complete sentences. This is an issue that requires further study and must take into account indirect speech acts (Perrault et al. 1980). The discourse stack accessed by our ellipsis processor (steps 2 and 14 in Figure 1) represents the system's current beliefs about discourse expectations, but our implementation does not include actually constructing it from the dialog preceding the elliptical fragment. Our intent in this section is to motivate the contents of the discourse stack immediately preceding an elliptical fragment.

For example, when IP asks a question of IS, IS is first expected to accept the question, either explicitly or implicitly, and then answer the question. Similarly, when IP answers a question posed by IS, IS is expected to accept the response, either explicitly or implicitly, and then pursue other discourse goals. This reasoning leads to the stack processing (SP) rules.

SP1: If IP asks a question of IS with the discourse goal of seeking information, Answer-Question and Accept Question are pushed onto the discourse stack.

SP2: When IP answers a question posed by IS, Accept-Response is pushed onto the discourse stack.

When IS pursues a discourse goal, such as seeking information or clarification, one's expectation is that IS will continue pursuing this discourse goal with subsequent utterances. This leads to the stack processing rule

SP3: When IS actively pursues a discourse goal, the discourse goal is pushed onto the discourse stack.

Although the strongest expectations are that IS will pursue a goal suggested by the top element of the discourse stack, this expectation can be passed over, at which point it no longer suggests expectations for utterances. This produces the stack popping (SPP) rule

SP4: When IS's utterance does not pursue a goal suggested by the top entry on the discourse stack, this entry is popped from the stack.

Other stack processing rules are similarly motivated and formulated. If IP asks or answers a question, SP1 or SP2 respectively apply. If IS makes an utterance, then rule SP4 is applied until it fails, at which point rule SP3 is applied.

7.3 SUGGESTED DISCOURSE GOALS

We have argued that an ongoing dialog establishes certain discourse expectations for IS and have represented these expectations as a discourse stack. We have further claimed that elliptical fragments do not explicitly indicate the discourse goal being pursued, and therefore IS must intend that elliptical fragments be interpreted according to mutually believed expectations. From our analysis of naturally occurring dialogs, we have formulated discourse expectation rules indicating how discourse expectations suggest the identity of IS's discourse goal. Associated with each discourse expectation is a rule that suggests a set of one or more discourse goals that IS might pursue and the order in which they should be considered.

Suppose that IP has posed a question to IS. Then the strongest discourse expectation is that IS will understand and accept the posed question or work toward understanding and accepting it. In order to understand a question, the question must have been satisfactorily transmitted. If not, IS will attempt to confirm those components he believes may have been miscommunicated; thus the first suggested discourse goal for IS under this discourse expectation is Seek-Confirm. Once IS believes the question has been properly transmitted, he must be able to satisfactorily identify the referents of the entities mentioned in the question. If he cannot do so, he may attempt to obtain further identification of an entity; thus the second suggested discourse goal for IS

under this discourse expectation is Seek-Identify. Once the components of the question are understood, IS must believe he fully comprehends what is being asked. If he does not, he must obtain clarification of the question before he can answer it; thus the third suggested discourse goal for IS under the discourse expectation of accepting the question is Seek-Clarify-Question. Once the question is understood, IS may refuse to accept it by expressing surprise at an inference drawn from the question or at the fact that the question was asked at all. Thus the fourth suggested discourse goal under the discourse expectation of accepting the question is Express-Surprise-Question. This leads to the following discourse expectation (DE) rule:

Rule DE1: The discourse expectation Accept-Question suggests the following ordered set of discourse goals for IS:

1. Seek-Confirm
2. Seek-Identify
3. Seek-Clarify-Question
4. Express-Surprise-Question

In a cooperative dialog, once IS has understood and accepted a question, IP expects him to answer it and is therefore on the lookout for an answer. Therefore first preference is given to interpretations that accomplish this goal; thus the first discourse goals suggested for IS in this situation are those that answer the question, directly or indirectly. If IS is unsure of the best answer to the question, he may suggest one or more possible answers and request IP's help in evaluating them; thus the second discourse goal suggested for IS is Suggest-Answer-Question. In a cooperative dialog, we expect IS to answer questions posed to him. If he does not do so, he must not have sufficient knowledge to formulate a good answer; in this case, we expect IS to work toward being able to answer the question by gathering whatever extra information he needs. Thus the third discourse goal suggested for IS under the discourse expectation of answering the question is Obtain-Information. This analysis leads to the following discourse expectation rule:

Rule DE2: The discourse expectation Answer Question suggests the following partially ordered set of discourse goals for IS:

1. Answer-Question
 - Answer-Question-With-Restrictions
 - Answer-Question-Suggest-Alternative
2. Suggest-Answer-Question
3. Obtain-Information

If IS has been informing IP about certain aspects of his underlying task-related plan, our strongest expectation is that he will elaborate on the plan by providing further detail. However, in an information seeking dialog, we anticipate that IS is providing this knowledge as background information in order to help IP formulate cooperative responses to subsequent questions. Thus our second expectation is that IS will request information

about the domain in order to achieve his objective of constructing a plan for his underlying task. This leads to the following discourse expectation rule:

Rule DE3: The discourse expectation Provide-For-Assimilation suggests the following ordered set of discourse goals for IS:

1. Provide-For-Assimilation
2. Obtain-Information

Other discourse expectation rules are similarly motivated and formulated.

As described earlier, discourse expectations are represented in the discourse stack. These discourse expectations suggest anticipated behavior for IS and are analyzed in turn until a coherent interpretation for the elliptical fragment is identified (Loop A in Figure 1). The top element of the discourse stack contains the strongest discourse expectation at the current point in the dialog. The rule associated with this discourse expectation is used to suggest discourse goals that IS might pursue (steps 2 and 3 in Figure 1). Each of these suggested discourse goals is tried in turn (Loop B in Figure 1). If it is determined that IS is not pursuing any of the discourse goals suggested by a discourse expectation, then that discourse expectation is popped from the discourse stack (step 14 in Figure 1) and the process is repeated (Loop A in Figure 1) using the rule associated with the new discourse expectation now residing on top of the stack to suggest discourse goals for IS. Thus, for example, when IP asks IS a question, the discourse expectations Answer-Question and Accept-Question are successively pushed onto the discourse stack. If IS then uses an elliptical fragment, the discourse expectation Accept-Question will first be used to suggest discourse goals that IS might be pursuing. However, if an interpretation of the fragment accomplishing one of these suggested discourse goals cannot be constructed, then this discourse expectation will be popped from the discourse stack; this indicates that IS has understood and implicitly accepted the question asked by IP. The top element of the stack will now be Answer-Question and it will be used to suggest discourse goals that IS might be pursuing via the fragment. (Note that although discourse expectations can be popped from the discourse stack, a cooperative dialog participant will never pass over the discourse goal of answering a question. In this research, we have assumed that the participants are always cooperative.)

7.4 DISCOURSE GOAL RULES

We have claimed that processing of elliptical fragments should be controlled by the discourse expectations in effect at that point in the dialog. Discourse expectations suggest discourse goals that IS might be expected to pursue. If an elliptical fragment can be interpreted as pursuing a suggested discourse goal, then this interpretation of the fragment should be regarded as the one intended by IS; otherwise IS would have overridden

these expectations, since he is aware that they will be used in understanding his utterances.

Associated with each discourse goal is one or more discourse goal rules (step 6 in Figure 1): each rule applies factual and processing knowledge to analyze a fragment and determine whether it can be understood as pursuing the discourse goal with which the rule is associated (Loop C in Figure 1). We view fragments as highlighting terms or propositions in the underlying task-related plan; recognition of the highlighted aspect is described in the next section. The discourse goal rules determine whether IP believes that it is mutually believed that IS might be pursuing the discourse goal under consideration and whether the highlighted aspect of the plan provides sufficient information for recognizing the suggested discourse goal (step 10 in Figure 1). If so, then the elliptical fragment should be understood as directed toward accomplishing this suggested discourse goal and a related interpretation produced (steps 11–13 in Figure 1).

In order to demonstrate how discourse goal rules facilitate recognition of the discourse goal that IS is executing, let us examine how IS might seek identification of an entity appearing in IP's utterance. Figure 2

<u>Seek-Identify-1(IS, entity1)</u>	
Applicability	$\neg \text{Know}(\text{IS}, \text{Referent}(\text{entity1}))$
Conditions:	$\text{Want}(\text{IS}, \text{Know}(\text{IS}, \text{Referent}(\text{entity1})))$
Body:	$\text{Request}(\text{IS}, \text{IP}, \text{Inform}(\text{IP}, \text{IS}, \text{Referent}(\text{entity1})))$
Effect:	$\text{Seek-Identify}(\text{IS}, \text{entity1})$
<u>Seek-Identify-2(IS, entity1, entity2)</u>	
Applicability	$\neg \text{Know}(\text{IS}, \text{Referent}(\text{entity1}))$
Conditions:	$\text{Want}(\text{IS}, \text{Know}(\text{IS}, \text{Referent}(\text{entity1})))$ $\text{Know}(\text{IS}, \text{Referent}(\text{entity2}))$ $\neg \text{Believe}(\text{IS}, \neg \text{Same-Referents}(\text{entity1}, \text{entity2}))$
Body:	$\text{Request}(\text{IS}, \text{IP}, \text{Inform}(\text{IP}, \text{IS}, \text{Same-Referents}(\text{entity1}, \text{entity2})))$
Effect:	$\text{Seek-Identify}(\text{IS}, \text{entity1})$

Figure 2 Two Plans for Seeking Identification of an Entity.

presents two plans, Seek-Identify-1 and Seek-Identify-2, whose primary effect is the same, namely to seek identification of an entity. These two plans differ in that in the body of the first plan, IS directly requests further identification of the entity in question, whereas in the second plan, IS attempts to determine the entity by asking if it is the same as an entity with which he is familiar. Now let us consider how an elliptical fragment might be recognized as pursuing one of these plans. Suppose we have the following dialog:

IS: "What CS courses are being offered during the summer?"
IP: "CS105 and CS461."

IS: "CS461?"

Suppose further that Seek-Identify has been suggested as a discourse goal for IS. If IS intends IP to recognize Seek-Identify-1 as the discourse plan that he is pursuing, then he must indicate entity1, the entity whose identification he is requesting. The elliptical fragment "CS461?" may be understood as doing this. The applicability conditions for Seek-Identify-1 specify constraints that must be satisfied in order for IS to pursue this plan; only if IP believes that it is mutually believed that these constraints are satisfied can IP infer that IS intends him to recognize that he is pursuing this discourse goal. In the above dialog, IP believes that it is mutually believed that IS wants to know the referent of CS461 since CS461 was mentioned in IP's previous response; it is also reasonable for IP to believe that IS may not know the referent of CS461.

Let us examine this latter belief in greater detail. One might ask why IP did not provide an extended description of CS461 in her previous response, if she believes that IS cannot satisfactorily identify the course from its department and number. This is not an issue with which this paper is concerned, but we will suggest one possibility. In constructing a description, IP uses knowledge about the listener (Appelt 1985, Goodman 1986) to produce a concise reference that she believes will be acceptable. Sometimes she will be quite certain that the listener will be able to identify the referent of her description—for example, perhaps the same description has been used successfully before. At other times, IP will be uncertain whether the description is sufficiently detailed; she could go to the extreme and include every detail at her disposal, but this may confuse the listener and hamper identification (Goodman 1986). Therefore, if IP does not know that IS can identify the referent of an entity from her description, she must wait for IS's response to this description to determine its success. In the case of the above dialog, IS does not implicitly or explicitly accept the description (he could do this by passing over the discourse goal of Seek-Identify) and IP is left with the belief that her description was unsatisfactory.

In contrast, consider the following variation of the above dialog:

IS: "What courses is Dr. Jones teaching next fall?"

IP: "CS440 and CS461."

IS: "What days is CS461 offered?"

IP: "Monday evenings."

IS: "What courses are being offered during the summer?"

IP: "CS105 and CS461."

IS: "CS461?"

It is unlikely that IS's last utterance would be interpreted as seeking identification of CS461. This is because IS implicitly accepted IP's first response, implying that he understood the description "CS461", and

subsequently used it himself. This leads IP to believe that IS can satisfactorily identify the referent of the term "CS461". Therefore, if IS really does want to pursue this discourse goal, he must override IP's belief with an utterance such as

"I haven't heard of CS461. What course is it?"

This leads to the discourse goal (DG) rule:

Rule DG-Seek-Identify-1: Check that the following conditions are satisfied:

1. IS's elliptical fragment terminates in a "?"
2. The fragment highlights a component of IS's underlying task-related plan and matches a description used in the utterance by IP that is closest to the top of the discourse stack.
3. It is mutually believed that IS might not know the referent of this description.

If these conditions are satisfied, then interpret the fragment as seeking further identification of the highlighted plan component.

The first condition in the above rule checks that IS is making a request (the body of the plan for Seek-Identify-1 in Figure 2), the second condition checks that IS has provided a component of his plan (which might be the entity whose referent is the object of the request in the Seek-Identify-1 plan in Figure 2) and that it is reasonable to believe that he wants to know its referent (the second applicability condition in the Seek-Identify-1 plan in Figure 2), and the third condition checks that IS does not already know what this referent is (the first applicability condition in the Seek-Identify-1 plan in Figure 2). A discourse goal rule is only invoked when discourse expectations suggest that IS might be pursuing its associated discourse goal. Therefore, if the conditions of an invoked rule are satisfied, then IS has provided all the information necessary for IP to believe that IS is pursuing that goal and IP is justified in interpreting the elliptical fragment accordingly.

Similarly, if IS intends IP to recognize Seek-Identify-2 as the discourse plan that he is pursuing, then its applicability conditions must be satisfied and IS must indicate both entity1 (the entity whose identification he is requesting) and entity2 (the entity with which he is familiar). Entity2 must be provided explicitly by the elliptical fragment, since it is new information not currently focused on in the dialog, and entity1 implicitly by the plan element with which the fragment is associated. This leads to the following discourse goal rule:

Rule DG-Seek-Identify-2: Check that the following conditions are satisfied:

1. IS's elliptical fragment terminates in a "?"
2. The fragment highlights a component of IS's underlying task-related plan and this component is referenced by a description D in the utterance by IP that is closest to the top of the discourse stack.

3. It is mutually believed that IS might not know the referent of the description D.
4. It is mutually believed that IS does not believe that the fragment and the description D refer to different entities⁶.

If these conditions are satisfied, then interpret the fragment as seeking further identification of the highlighted plan component. In particular, IS's fragment describes an entity with which he is familiar and asks whether this entity and the highlighted plan component are the same.

One might suggest using a general plan recognition strategy to recognize discourse goals as well as domain goals. However, such a strategy has not been adequately worked out even for complete utterances, although Litman's work (Litman and Allen 1987) is a major step in this direction. Furthermore, as we argued earlier, elliptical fragments require special treatment. Our discourse goal rules are compilations of the kind of reasoning that must be done on discourse plans to ascertain if an elliptical fragment is intended to be recognized as pursuing a suggested discourse goal.

The following are four other sample discourse goal rules. As mentioned earlier and discussed in detail in the next section, an elliptical fragment is viewed as highlighting a term T or proposition P in IS's plan, possibly within the context of an additional set of propositions, which we will refer to in the following rules as CONTEXT PROPS. For example, the fragment "With Dr. Smith?" might highlight the proposition

P = Teach(SMITH,ss:&SECTIONS),

but the plan inferred from the preceding dialog might indicate that P should be interpreted within the context of the additional propositions

Is Section Of(ss:&SECTIONS, CS200)

Section Offered(ss:&SECTIONS,
NEXT-SEMESTER)

If IS's discourse goal were identified as Obtain Information, the resultant interpretation would be that IS wants to know whether Dr. Smith is teaching a section of CS200 next semester, not merely whether Dr. Smith is *ever* teaching a section of *any* course. The examples in Section 9 will illustrate the use of the following rules.

Rule DG-Express-Surprise-Obtain-Corroboation-1:

Check that the following conditions are satisfied:

1. IS's elliptical fragment terminates in a "?"
2. The fragment highlights a proposition P in the context of a conjunction of propositions CONTEXT-PROPS in IS's underlying task-related plan.
3. It is mutually believed that IS already knows IP's beliefs about the truth of the proposition $(P \wedge \text{CONTEXT-PROPS})$.

If these conditions are satisfied, then interpret the fragment as expressing surprise at IP's response and

seeking corroboration of it. In particular, IS is surprised at the known truth value of $(P \wedge \text{CONTEXT-PROPS})$ in light of the new information provided by IP's response and the proposition P highlighted by IS's fragment.

Rule DG-Obtain-Information: Check that the following conditions are satisfied:

1. IS's elliptical fragment terminates with a "?"
2. The fragment highlights a term T or proposition P in the context of a conjunction of propositions CONTEXT-PROPS in IS's underlying task-related plan.
3. It is mutually believed that IS does not know the value of the highlighted term or proposition.

If these conditions are satisfied, then interpret the fragment as requesting information about the value of the highlighted term T or proposition P within the context of the propositions in CONTEXT PROPS.

Rule DG-Provide-For-Assimilation: Check that the following conditions are satisfied:

1. IS's elliptical fragment terminates with a "."
2. The fragment highlights a term T or proposition P in the context of a conjunction of propositions CONTEXT-PROPS in IS's underlying task-related plan.

If these conditions are satisfied, then interpret the fragment as specifying that the user-specified term replace the term with which it is associated in the plan or that the proposition be satisfied as part of the plan being constructed.

Rule DG-Answer-Question-With-Restrictions-1:

Check that the following conditions are satisfied:

1. IS's elliptical fragment terminates with a "."
2. The last unanswered question asked by IP (the one at the top of the discourse stack) was a Yes-No question.
3. The elliptical fragment highlights one or more propositions P in the context of a conjunction of propositions CONTEXT-PROPS in IS's underlying task-related plan.

If these conditions are satisfied, then interpret the fragment as answering *yes*, with the restriction that the proposition $(P \wedge \text{CONTEXT-PROPS})$ be satisfied in the underlying plan.

8 PLAN ANALYSIS

8.1 REPRESENTATION OF PLANS

We view a plan as the means by which an agent can accomplish a nonprimitive, task-related goal. Plans are represented, using an extended STRIPS formalism (Fikes and Nilsson 1971; Allen et al. 1980), as structures containing applicability conditions, preconditions, a plan body, and effects. The plan body is a decomposition of the plan's goal into a set of simpler subgoals, each of which can be accomplished by a primitive action or itself has an associated domain plan. Arguments in plans are either constants, represented as uppercase

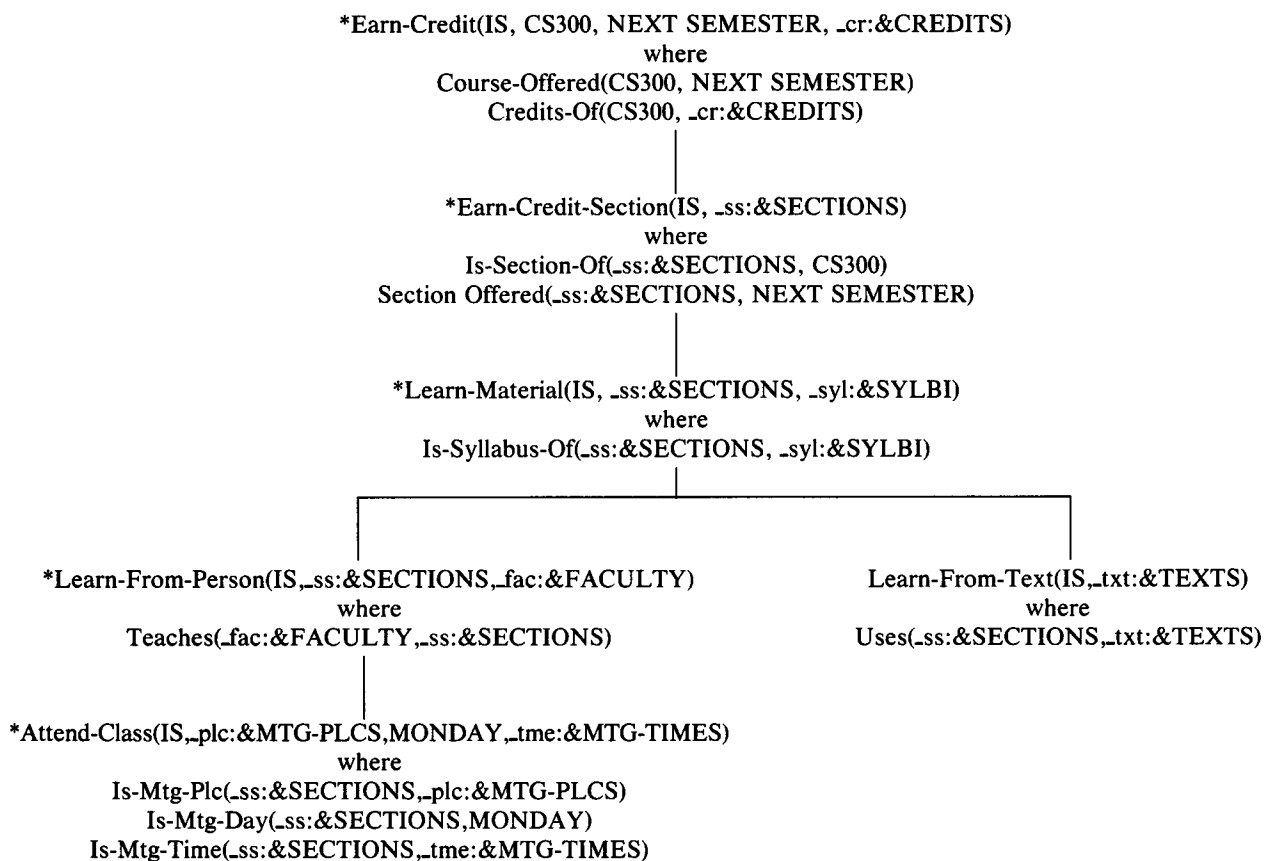


Figure 3 A Portion of an Expanded Context Model.

strings, or typed variables, represented as lowercase strings preceded by the character “_” and followed by the characters “:&” and an uppercase string giving the variable’s type. These plans are hierarchical, since many of the subgoals in a plan are nonprimitive and therefore have associated plans that may be substituted for them. Thus a plan can be expanded to any desired degree of detail, by repeatedly replacing subgoals with associated plans, which themselves contain constituent subgoals.

We use a tree structure called a **context model** to represent the task-related plan inferred for IS from the preceding dialog. Each node in this tree represents a goal that IS has investigated achieving and, except for the root, is a descendant of a higher-level goal whose associated plan contains the goal represented by the child node. For example, in the context model shown in Figure 3, IS is interested in achieving the subgoal of learning the material in a section of CS300 and has investigated learning from the instructor and learning from the text. Since only those subgoals that have been investigated as part of a plan for achieving a goal are included in the context model, a goal node can have different children in two context models derived from different dialogs. One node in the context model is marked as the current focus of attention and indicates that aspect of the task on which IS’s attention is

currently centered. The path from the root of the context model to the current focus of attention is called the **active path** and represents the global context, or sequence of progressively lower-level goals that led to the subgoal under consideration by IS. Carberry (1983, 1988) describes how the context model is dynamically constructed from an ongoing dialog.

8.2 ASSOCIATION OF FRAGMENTS

We view elliptical fragments as highlighting, or calling attention to, a term or proposition in IS’s underlying task-related plan. Given an elliptical fragment, the task-related plan inferred for IS must be analyzed to determine, based on the current focus of attention, the particular aspect of the plan highlighted by IS’s fragment. This paper will discuss three classes of elliptical fragments; a description of how other fragments are associated with plan elements is provided in Carberry (1985).

We will call the matching of fragments with plan elements **association**. A constant fragment can only associate with terms whose semantic type is the same or a superset of the semantic type of the constant. Furthermore, each term in a plan has a limited set of valid instantiations. A constant associates with a term only if the system’s beliefs indicate that IS might believe that

the uttered constant is one of the term's valid instantiations. For example, if a plan contains the proposition

Starting-Date(AI-CONF, _date:&DATES)

the elliptical fragment "February 2?" will associate with the date term in this proposition only if the system believes IS might believe that the starting date for the AI conference is in February. Recourse to such a belief model is necessary in order to allow for Yes-No questions to which the answer is "No" and yet eliminate potential associations which a human listener would recognize as unlikely. Sidner (1981) employs a similar strategy in her work on anaphora resolution. A co-specifier proposed by the focusing rules must be confirmed by an inference machine; if any contradictions are detected, other co-specifiers are suggested.)

A propositional fragment can be of two types. The first contains a proposition whose name is the same as the name of a proposition in the plan domain. The second type is a more general propositional fragment, which cannot be associated with a specific plan-based proposition until after analyzing the relevant propositions appearing in IS's plan. The semantic representations of the utterances

"Taught by Dr. Smith?"

"With Dr. Smith?"

would produce respectively the type 1 and type 2 propositions

Teaches(SMITH, _ss:&SECTIONS)

Genpred(SMITH)

The latter indicates that the name of the specific plan proposition is as yet unknown, but that one of its parameters must associate with the constant Smith.

A proposition of the first type associates with a proposition of the same name if the parameters of the propositions associate. A proposition of the second type associates with any proposition whose parameters include terms associating with the known parameters of the propositional fragment.

The semantic representation of a term such as "The meeting time?" contains a variable term

_tme:&MTG-TIMES

Such a term associates with terms of the same semantic type in IS's plan. Note that the existing plan may contain constant instantiations in place of former variables. A term fragment still associates with such constant terms.

8.3 RETAINING ESTABLISHED CONTEXT

It appears that humans retain as much of the established context as possible in interpreting intersentential ellipsis. Carbonell (1983) demonstrated this phenomenon in an informal poll in which users were found to interpret the fragment in the following dialog as retaining the fixed media specification.

IS: "What is the size of the 3 largest single port fixed media disks?"

IP:

IS: "Disks with two ports?"

We have noted the same phenomenon in other domains. Thus, when an elliptical fragment is associated with a component of the task-related plan, the context established by the preceding dialog should be used to replace information deleted from this streamlined, fragmentary utterance.

As mentioned earlier, we use a context tree to represent the underlying task-related plan inferred for IS at this point in the dialog. The set of nodes along the path from the root of the context tree to the current focus of attention form a stack of goals and associated plans, the topmost of which is the most recently considered subgoal in the current focused plan; each of these goals is part of the plan associated with the goal immediately beneath it in this stack. These **active nodes** represent the established global context, and the propositions in these nodes can have the effect of restricting the instantiation of variables in their ancestor nodes in the context tree. For example, in Figure, IS's most recently considered subgoal is attending class on Monday; this restricts the sections that he is currently considering taking to those which meet on Monday, thus limiting the possible instantiations of the variable _ss in the Learn From Person, Learn Material, and Earn Credit-Section subgoals higher up on the active path. These restrictions represent part of the context within which the elliptical fragment occurs and, unless overridden by information in the utterance itself, represent information that is missing from the sentence fragment but is understood by the speaker. Thus, if we use the propositions appearing along the active path, we can retain the restrictions on the variables in ancestor nodes when the focus of attention pops back to them. On the other hand, the nodes along the paths from the root of the context tree to the nodes at which fragment elements associate with plan elements represent the new context derived from IS's fragmentary utterance.

For example, consider the following dialog:

IS: "What days is CS300 taught next semester?"

IP: "Two sections of CS300 meet on Monday and one section meets on Tuesday."

IS: "What time do the sections on Monday meet?"

IP: "One section meets on Monday at 4 pm and another section meets on Monday at 7 pm"

IS: "The texts?"

A portion of the expansion of IS's inferred task-related plan is shown in Figure 3. Active nodes are marked by asterisks. The semantic representation of the fragment "The texts?" will contain the variable term _book:&TEXTS. This term associates with the term _txt:&TEXTS appearing at the node for the goal

Learn-From-Text(IS, _txt:&TEXTS)

where

Uses(_ss:&SECTIONS, _txt:&TEXTS)

The propositions along the active path are

Course-Offered(CS300, NEXT SEMESTER)

Credits-Of(CS300, _cr:&CREDITS)

Is-Section-Of(_ss:&SECTIONS, CS300)

Section-Offered(_ss:&SECTIONS,
NEXT-SEMESTER)

Is-Syllabus-Of(_ss:&SECTIONS, _syl:&SYLBI)

Teaches(_fac:&FACULTY, _ss:&SECTIONS)

Is-Mtg-Plc(_ss:&SECTIONS, _plc:&MTG-PLCS)

Is-Mtg-Day(_ss:&SECTIONS, MONDAY)

Is-Mtg-Time(_ss:&SECTIONS,
_tme:&MTG-TIMES)

These propositions maintain the established context—namely, that we are talking about the sections of CS300 next semester that meet on Monday. The path from the root of the context tree to the node at which the elliptical fragment associates with a term in the plan produces the additional proposition Uses(_ss:&SECTIONS, _book:&TEXTS). Thus the fragment highlights the term _book:&TEXTS within the context of the conjunction of this proposition and the nine propositions appearing in the active nodes. (These 10 propositions are referred to as CONTEXT-PROPS in the discourse goal rules of Section 7.4.) The semantics of this interpretation is that IS is drawing attention to the term _book:&TEXTS such that the conjunction of all 10 propositions is satisfied—namely, textbooks used in sections of CS300 that meet on Monday next semester.

8.4 IDENTIFYING THE INTENDED ASSOCIATION

Within IS's task-related plan, there often are multiple components with which a fragment might associate. Therefore it is necessary to identify that aspect of the plan that IS intended to draw attention to by using the elliptical fragment. Grosz has claimed that communication can be successful only if both speaker and listener are focused on the same subset of knowledge and if the listener recognizes any shifts in focus by the speaker. Grosz (1977) formulated rules for recognizing focus shifts in apprentice-expert task dialogs. McKeown (1985) expanded on focus rules proposed by Sidner (1981) to explain how speakers should organize their utterances when faced with a choice of topic.

We have investigated information-seeking dialogs in which the information-seeker is attempting to construct a plan for a task to be executed at some time in the future. These dialogs are not constrained by the order of execution of the steps of the underlying task, as in apprentice-expert dialogs. However, such dialogs do

exhibit structure in that participants typically move around in predictable ways to discuss different aspects of the plan being constructed. This structure has led to the development of focusing rules that predict how a speaker may shift his focus of attention within the underlying task-related plan (Carberry 1983, 1988).

An elliptical fragment by itself often contains very little information and the fragment is likely to be misinterpreted if its relationship to the dialog context (represented, in our case, by the discourse stack and the context model) is not correctly identified. But since discrepancies in how participants view their knowledge may exist, speakers using elliptical fragments cannot rely on finely grained plan structures for correct interpretation. Thus small shifts in focus of attention in the context model are not significant in processing elliptical fragments.

Although a shift in focus occurs every time one moves from a goal to a subgoal in the goal's associated plan, some of these moves may represent relatively small shifts in attention while others may be quite large. We employ the notion of focus domains in order to group together goals that appear to be at approximately the same level of focus when a plan is explicitly focused. Moving from one goal to another in the same focus domain will be considered a smaller shift in attention than moving from a goal in one focus domain to one in a different focus domain.

In a context model, a child node will often be in the same focus domain as its parent. However, if investigating how to achieve a goal represents a significant shift in focus from the plan that contains it, the goal is marked in the plan library as introducing a new focus domain. This means that all of the goal's children (representing subgoals in its plan) will be part of a different focus domain than the goal itself—i.e., a move from discussing the goal to considering how one would go about achieving it represents a large shift in attention. In Figure 4, the subgoals preceded by a "B" are in a different focus domain than those preceded by an "A". The goals of taking a particular section of a course, learning the material of a course, and learning from a particular teacher all reside in the same focus domain within the expanded plan for earning credit in a course (and therefore are in focus to approximately the same degree). The goal of going to the cashier's office to pay one's tuition also appears within this expanded plan; however, it is part of a different focus domain, introduced by the goal of paying tuition, since it does not come to mind nearly so readily when one thinks about taking a course.⁷ The goal Pay-Tuition(IS, _cr:&CREDITS) is marked in Figure 4 as A/B, indicating that it is in focus domain A but that it introduces focus domain B (the goals in its associated plan are in focus domain B).

The focusing rules use mutual beliefs about knowledge currently focused on by the dialog participants and expectations about probable shifts in focus to rank

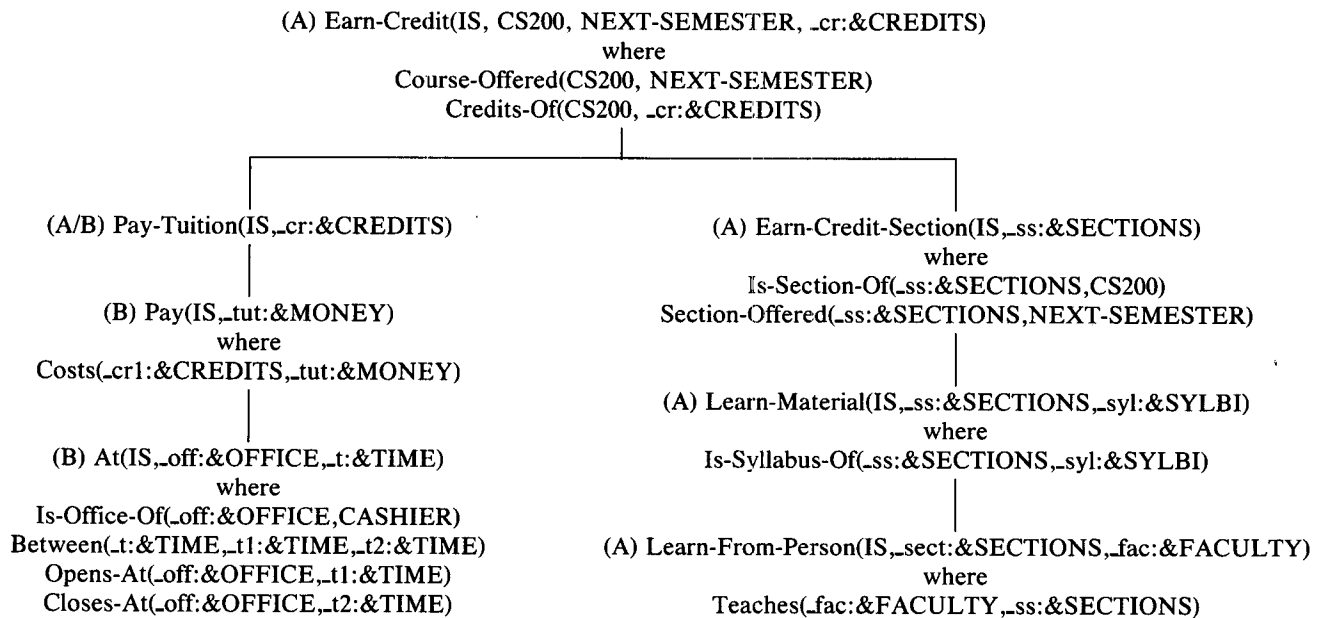


Figure 4 Part of a Context Model Illustrating Two Focus Domains.

alternative associations of elliptical fragments with elements of IS's underlying task-related plan. The current focus domain contains those goals that are most highly focused within IS's underlying task-related plan at the current time; therefore interpretations relevant to these goals are preferred. For example, consider the following dialog segment:

IS: "Who is teaching CS440?"
 IP: "Dr. Smith."
 IS: "What textbook is being used?"
 IP: "There is no textbook. However, a set of papers and notes have been assembled into a course manual."
 IS: "Does the bookstore carry the course manual?"
 IP: "Yes."
 IS: "Any faculty discounts?"

Faculty discounts might be available on both courses taken by faculty and books purchased at the bookstore. But humans interpret IS's elliptical fragment as asking whether there is a discount for faculty purchases of the course manual at the bookstore. This is explained by noting that the goal of purchasing a text is part of the current focus domain at the time the elliptical fragment is uttered, whereas the goal of paying for the course is not, and interpretations within the current focus domain or its descendants are preferred over other interpretations in the overall plan. This leads to the rule:

Rule-Current-Focus-Domain: Associations with elements in the current focus domain, or within expansions of the plans associated with goals in the current focus domain, should be preferred.

Although the majority of elliptical fragments require little shift in focus, fragments requiring significant focus

shift do occur and are generally interpreted correctly by human dialog participants. Consider for example the following dialog:

IS: "Who is teaching CS440?"
 IP: "Dr. Smith."
 IS: "What textbook is being used?"
 IP: "There is no textbook. However, a set of notes and papers have been assembled into a course manual."
 IS: "Does the bookstore carry the course manual?"
 IP: "Yes."
 IS: "3 credits?"

The fragment "3 credits?" should be interpreted as asking whether CS440 is a 3 credit hour course. Purchasing course materials appears as a subtask in the plan for taking CS440, and interpreting IS's elliptical fragment requires shifting from the focus domain associated with this subtask to the focus domain associated with taking CS440. This is accounted for by the expectation that IS will eventually return to the higher level plans whose expansion led to the most recently considered goal, leading to the following focus rule:

Rule-Higher-Plans: Associations with elements in the focus domains of higher level plans whose expansion led to the current focus domain should be preferred over other associations with elements that do not appear in an expansion of a goal in the current focused domain.

Other focusing rules are similarly motivated.

Although un signaled focus shifts may sometimes seem awkward, human dialog participants appear to recognize their intended meaning as long as the utter-

ances are specific enough to avoid confusion. As an example, consider the following naturally occurring dialog that took place in early spring.

Speaker 1: "Have you received the return receipt for your conference paper?"

Speaker 2: "Yes, it came yesterday."

Speaker 1: "August the ___?"

The fragment "August the ___?" was correctly interpreted by Speaker 2 as a request for the date that Speaker 2 would be attending the conference to which the paper had been submitted.

In a few instances, an elliptical fragment may associate with two alternative elements of IS's underlying task-related plan, with the focusing heuristics rating both equally as the association intended by the speaker. In this case, the utterance will be interpreted as addressing both aspects of the speaker's plan. This is discussed further in Carberry (1985).

9 IMPLEMENTATION AND EXAMPLES

A prototype system demonstrating this pragmatics-based framework for processing intersentential ellipsis has been implemented for a domain consisting of the courses, policies, and requirements for students at a university. The system is presented with a semantic representation of an elliptical fragment, a context tree representing the system's beliefs about the information seeker's underlying task-related plan, a belief model representing the system's beliefs about the information seeker's beliefs, and an initial discourse stack containing the system's beliefs about the information seeker's expected discourse behavior. The fragment's semantic representation gives the structure of the fragment as discussed in Section 8.2 and its terminating punctuation. The context tree is inferred and constructed from the preceding dialog using the incremental plan recognition algorithm implemented in the TRACK system and described in Carberry (1988). Our belief model is very primitive and is only intended to facilitate illustration of the ellipsis processor. The belief model contains propositions that the system believes the information-seeker believes are satisfied. These propositions may be fully or partially instantiated; in the latter case, they represent the belief that the information seeker believes the proposition can be satisfied by a constant selected from the designated type for the argument. So, for example, the system's belief that the information-seeker believes that the starting date of the AI conference is in February would be represented in the belief model by the proposition

Starting Date(AI-CONF,
_date:&FEBRUARY-DATES)

where FEBRUARY-DATES is a subclass of DATES. The belief model also contains other beliefs, such as the objects and types that the information-seeker knows

about and the information-seeker's beliefs about the types of constants and about type generalization. Since our research has not addressed issues such as how beliefs are acquired during a dialog (Kass et al. 1987), how they can be efficiently represented and updated in a user model (Finin et al. 1986), or how discourse goals are recognized from complete sentences, our implementation does not include construction of the belief model or the discourse stack from the dialog preceding the elliptical fragment.

The following examples illustrate our ellipsis resolution strategy.⁸ All rules referenced in the first example have been described in the preceding sections. The other examples illustrate a variety of discourse situations but, due to space limitations, are presented in less detail than the first example.

9.1 EXAMPLE 1: OBTAINING INFORMATION

IS: "I want to register for a course for next semester. But I missed pre-registration. The cost?"

In this example, the first two utterances establish a plan context of taking a course, with attention directed to the subtask of late registering for it, and the elliptical fragment should be interpreted as the cost of registering late. Since proper understanding of the fragment relies on inferred knowledge about IS's underlying task, this example requires a plan-based interpretation.

IS's first two utterances provide information for the system to use in inferring IS's underlying task-related plan. An expansion of the context tree at this point in the dialog is shown in Figure 5. In the belief model, the only proposition about costs asserts that IS believes that the cost of any item is of type MONEY; therefore the system assumes that IS does not know the specific cost of individual items. The semantic representation of the elliptical fragment is the triple

_cost1:&MONEY Definite ?

specifying that the fragment is a definite noun phrase whose head noun is of type MONEY, followed by a question mark. Based on the dialog preceding the elliptical fragment, the discourse stack contains the entry Provide-For-Assimilation (discourse stack rule SP3 in Section 7.2). Its associated discourse expectation rule, Rule DE3, suggests that the strongest expectation is for IS to continue providing information to the system and that the next strongest expectation is for IS to seek information in order to construct a plan for the task motivating his interaction with the system. The discourse goal rule associated with Provide-For-Assimilation, Rule DG-Provide-For-Assimilation, fails to find an interpretation since the elliptical fragment terminates with a "?". Next, the discourse goal rule associated with Obtain Information, Rule DG-Obtain-Information, is invoked. It checks that the fragment terminates with

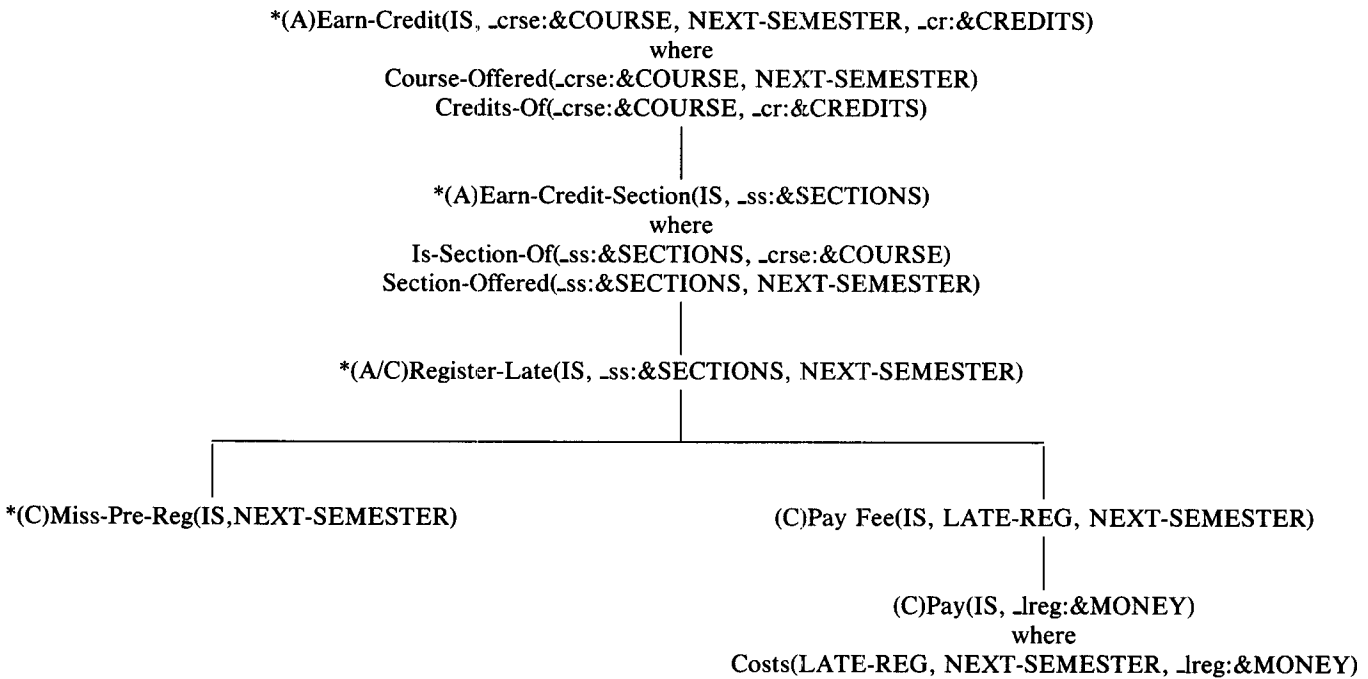


Figure 5 A Portion of the Expanded Context Tree for Example 1.

a period and calls the plan analysis component to associate the fragment with IS's inferred underlying task-related plan.

In the expanded context tree shown in Figure 5, the parenthesized letters preceding goals indicate each goal's focus domain, and nodes on the active path are marked by asterisks. Immediately prior to IS's elliptical fragment, the current focused plan is the plan associated with the goal of registering late, the most recently considered aspect of this plan is missing pre-registration, and the current focus domain contains the goals preceded by (C). IS's fragment associates with the term `_lreg:&MONEY` in IS's inferred plan, as well as with terms elsewhere in parts of the expanded plan not shown in Figure 5. However, none of the other terms appear in the current focus domain, and therefore the association of the fragment with `_lreg:&MONEY` is selected as most relevant to the current dialog context. The path from the root of the context model to the current focus of attention prior to the elliptical fragment contains the propositions

```

Course-Offered(_crse:&COURSE,
               NEXT SEMESTER)
Credits-Of(_crse:&COURSE, _cr:&CREDITS)
Is-Section-Of(_ss:&SECTIONS, _crse:&COURSE)
Section Offered(_ss:&SECTIONS,
                NEXT-SEMESTER)
  
```

indicating that the fragment should be interpreted within this context; the path to the node at which the fragment associates with a term in the plan provides the additional proposition

```

Costs(LATE-REG, NEXT SEMESTER,
      _cost1:&MONEY)
  
```

Since the belief model does not contain a proposition indicating that IS knows the value of the term `_cost1:&MONEY`, Rule DG-Obtain Information interprets the fragment as seeking information about the charge for late registration in order to formulate the task-related plan; in particular, IS is requesting the value of the term `_cost1:&MONEY` such that the conjunction of the following propositions (CONTEXT PROPS) is satisfied:

```

Course-Offered(_crse:&COURSE,
               NEXT-SEMESTER)
Credits-Of(_crse:&COURSE, _cr:&CREDITS)
Is-Section-Of(_ss:&SECTIONS, _crse:&COURSE)
Section-Offered(_ss:&SECTIONS,
                NEXT SEMESTER)
Costs(LATE-REG, NEXT SEMESTER,
      _cost1:&MONEY)
  
```

9.2 EXAMPLE 2: EXPRESSING SURPRISE

IS: "I want to take CS310 next semester.
Who is teaching it?"
IP: "Dr. Smith is teaching CS310 next semester."
IS: "What time does it meet?"
IP: "It meets at 8AM."
IS: "With Dr. Smith?"

This example illustrates an elliptical fragment that conveys IS's surprise at IP's response and seeks elaboration and corroboration of it. The semantic representation of the elliptical fragment is

Genpred(SMITH) Proposition ?

specifying that the fragment is followed by a question mark and is a general propositional fragment that must associate with a plan proposition that has Smith as an argument. The context tree immediately prior to the elliptical fragment is shown in Figure 6, with asterisks preceding nodes along the path from the root of the context model to the most recently considered goal. The belief model indicates that IS knows that Dr. Smith is teaching a section of CS310 next semester, since this was communicated and accepted prior to the fragment. The top two entries of the discourse stack are

Accept-Response
Obtain-Information

indicating that IS is first expected to accept IP's response to his previous question and then proceed to seek further information in order to expand his underlying task-related plan.

The discourse expectation rule associated with Accept-Response, similar to Rule DE1, suggests discourse goals of Seek-Confirm, Seek-Identify, and Express-Surprise-Obtain-Corroboration. The discourse goal rules associated with the first two of these discourse goals fail to produce an interpretation since the fragment does not associate with a term in IS's plan that was a component of IP's previous utterance. Rule DG-Express Surprise-Obtain Corroboration-1 checks that the fragment ends with a "?" and invokes plan

analysis on the context tree shown in Figure 6. It finds that IS's elliptical fragment associates with the proposition Teaches(_fac:&FACULTY, _ss:&SECTIONS) producing the instantiated proposition

P = Teaches(SMITH, _ss:&SECTIONS)

In order to retain the existing dialog context, this proposition must be interpreted within the context of the propositions appearing in nodes preceded by asterisks in Figure 6, namely

Course-Offered(CS310, NEXT SEMESTER)
Credits-Of(CS310, _cr:&CREDITS)
Is-Section-Of(_ss:&SECTIONS, CS310)
Section-Offered(_ss:&SECTIONS,
NEXT SEMESTER)
Is-Syllabus-Of(_ss:&SECTIONS, _syl:&SYLBI)
Is-Mtg-Plc(_ss:&SECTIONS, _plc:&MTG-PLCS)
Is-Mtg-Day(_ss:&SECTIONS,
_day:&MTG-DAYS)
Is-Mtg-Time(_ss:&SECTIONS,
_tme:&MTG-TIMES)

The discourse goal rules refer to the conjunction of the above eight propositions as CONTEXT PROPS. Since our belief model indicates that IS already knows the truth value of the proposition

(P \wedge CONTEXT-PROPS)

— namely, that Dr. Smith is teaching a section of CS310 next semester, Rule DG-Express-Surprise-Obtain-Cor-

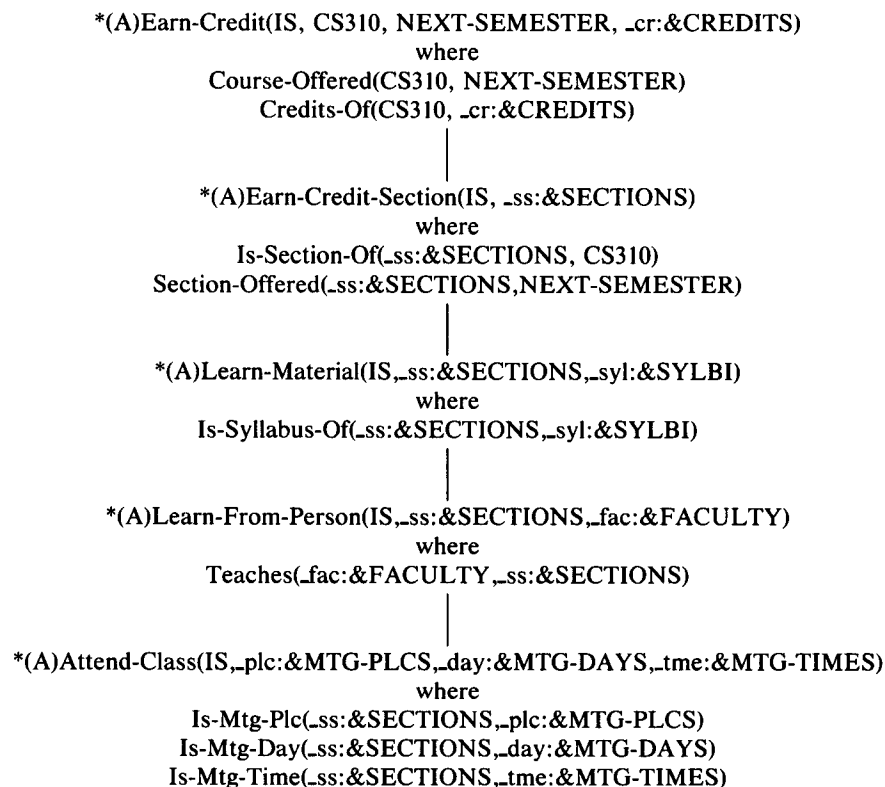


Figure 6 The Context Tree for Example 2.

roboration-1 identifies the elliptical fragment as expressing surprise at, and requesting corroboration of, IP's response. In particular, the system believes this surprise is a result of

1. the new information presented in IP's response, namely that 8 am is the value of the term `_tme:&MTG-TIMES` in the proposition `Is-Mtg-Time(_ss:&SECTIONS,_tme:&MTG-TIMES)` in `CONTEXT-PROPS`.
2. the aspect of the plan highlighted by IS's elliptical fragment, namely the proposition

`P = Teaches(SMITH,_ss:&SECTIONS)`

Precisely the reason why this data surprises IS would require an additional inference mechanism; perhaps Dr. Smith is known to be a notoriously late riser in the morning or perhaps Dr. Smith holds a full-time job elsewhere and only teaches occasional evening courses at the university.

9.3 EXAMPLE 3: POPPING THE DISCOURSE STACK

IP: "Do you want to take CS105 next semester?"

IS: "Who is teaching it?"

IP: "Dr. Ames and Dr. Wilks."

IS: "On Monday, Wednesday, Friday with Dr. Ames."

This example illustrates a situation in which multiple expectations must be popped from the discourse stack in processing the fragment. IP's first utterance establishes an expectation that IS will accept and answer the question asked by IP; therefore Answer Question and Accept-Question are pushed onto the discourse stack. As a result of IS's first utterance, the question is implicitly accepted, popping Accept Question from the discourse stack. IS's first utterance is a request for information in order to formulate an answer to IP's posed question, resulting in the expectation that IS will continue gathering such information; thus the expectation Obtain-Information is pushed onto the discourse stack. IP's response answers IS's question and causes the expectation Accept-Response to be pushed onto the discourse stack. Thus the discourse stack passed to the ellipsis processor contains the three entries

Accept-Response
Obtain-Information
Answer-Question

The context tree at this point in the dialog would be similar to that shown in Figure 6, except that CS105 would replace CS310 and only the first four nodes would be on the active path. The semantic representation of IS's elliptical fragment is the triple

[GenTimePred(MON-WED-FRI),
GenPred(AMES)] Proposition .

specifying that the fragment is a conjunction of two propositions, followed by a period.

The discourse expectation rule associated with the

top element of the discourse stack suggests the discourse goals Seek-Confirm, Seek-Identify, and Express Surprise-Obtain-Corroboration. The discourse goal rules associated with these discourse goals require that the fragment terminate with a question mark in order to produce an interpretation. Since this is not the case, the expectation Accept-Response is popped from the discourse stack. The new top element of the discourse stack is Obtain-Information. Its associated discourse expectation rule suggests the discourse goal of obtaining further information, since once IS pursues such a discourse goal, he is expected to continue pursuing it with subsequent utterances. Once again, since the fragment terminates in a ".", the associated discourse goal rule, Rule DG-Obtain-Information, fails to produce an interpretation, and the expectation Obtain-Information is popped from the discourse stack.

The discourse expectation rule associated with the expectation Answer-Question, Rule DE2, suggests, among others, the discourse goal of answering the question by providing restrictions on the underlying task-related plan. Rule DG-Answer-Question-With-Restrictions-1 invokes plan analysis on the context tree produced by the preceding dialog and finds that the propositions comprising the fragment associate with the plan propositions

`Is-Mtg-Day(_ss:&SECTIONS, _day:&MTG-DAYS)`
`Teaches(_fac:&FACULTY, _ss:&SECTIONS)`

resulting in the instantiated propositions

`Is-Mtg-Day(_ss:&SECTIONS, MON-WED-FRI)`
`Teaches(AMES, _ss:&SECTIONS)`

The nodes along the path from the root of the context model to the current focus of attention prior to the fragment indicate that the fragment should be interpreted within the context established by the additional propositions

`Course-Offered(CS105, NEXT-SEMESTER)`
`Credits-Of(CS105, _cr:&CREDITS)`
`Is-Section-Of(_ss:&SECTIONS, CS105)`
`Section-Offered(_ss:&SECTIONS,`
`NEXT-SEMESTER)`
`Is-Syllabus-Of(_ss:&SECTIONS, _syl:&SYLBI)`

Since the fragment terminates with a period and associates with plan propositions, discourse goal rule Rule DG-Answer-Question-With-Restrictions-1 interprets the question as answering affirmatively, with the restriction that the conjunction of the propositions

`Course-Offered(CS105, NEXT-SEMESTER)`
`Credits-Of(CS105, _cr:&CREDITS)`
`Is-Section-Of(_ss:&SECTIONS, CS105)`
`Section-Offered(_ss:&SECTIONS,`
`NEXT-SEMESTER)`
`Is-Syllabus-Of(_ss:&SECTIONS, _syl:&SYLBI)`
`Teaches(AMES, _ss:&SECTIONS)`

Is-Meeting-Day(,ss:&SECTIONS,
MON-WED-FRI)

be satisfied in IS's task-related plan—namely that the section of CS105 he enrolls in be taught by Dr. Ames and meet on Monday, Wednesday, and Friday next semester.

10 EXTENSIONS AND FUTURE WORK

The main limitation of this pragmatics-based framework appears to be in handling intersentential elliptical utterances such as the following:

IS: "Who is the teacher of CS200?"
IP: "Dr. Herd is the teacher of CS200."
IS: "CS263?"

Obviously, IS's elliptical fragment requests the teacher of CS263. Our model cannot currently handle such fragments. This limitation is caused by our mechanisms for retaining dialog context; they assume that IS constructs a plan for a task in a depth-first fashion, completing investigation of a plan for CS200 before moving on to investigate a plan for CS263. Since the teacher of CS200 has nothing to do with the plan for taking CS263, the mechanisms for retaining dialog context will fail to identify "teacher-of-CS263" as the information requested by IS.

One might argue that the elliptical fragment in the above dialog relies heavily on the syntactic representation of the preceding utterance and thus a syntactic strategy is required for interpretation. This may be true. However, dialogs such as the above really investigate task-related plans in a kind of "breadth-first" fashion; for example, in the above dialog, IS is analyzing the teachers of each course under consideration first, and will then move to considering other attributes of the courses. It appears that our plan-based framework can be extended to handle many such dialogs, by reasoning about how IS is constructing his task-related plan (Ramshaw 1989).

Litman (Litman et al. 1987) uses metaplans in her formalism and is able to handle the fragment in the above dialog. However, she views IS as changing the plan from consideration of CS200 to CS263, not as possibly considering several plans simultaneously. As a result, she cannot handle fragments such as IS's last utterance in the following dialog:

IS: "Who is teaching CS105?"
IP: "Dr. Smith and Dr. Jones are teaching sections of CS105."
IS: "Who is teaching CS106?"
IP: "Dr. Derr is teaching CS106."
IS: "The meeting time of Dr. Smith's section?"

It appears that successful handling of such fragments requires that we consider several potential plans in parallel and we are now considering mechanisms for doing so.

This research has considered only interpretation of sentence fragments. But it appears that the plan analyzer could be extended to provide information missing from a syntactically complete sentence. For example, suppose IS asked:

"Who is teaching CS300 on Monday next semester?"

and then subsequently asked:

"What textbooks will be used?"

The plan analyzer could identify both the element of the user's plan that associates with the incompletely specified proposition *UsesPred(Unknown,txt:&TEXTS)* appearing in the semantic representation of the user's utterance and the context within which the association occurs. This would enable the system to recognize that IS wants to know the texts *used by sections of CS300 that meet next semester on Monday*. This is similar to interpreting the elliptical fragment "The texts?" in Section 8.3. A related problem, that of understanding pragmatically ill-formed utterances (utterances that violate the system's model of the world), is addressed in Carberry (1988).

11 CONCLUSIONS

This research has investigated the use of pragmatic knowledge in understanding intersentential ellipsis occurring during an information seeking dialog in a task domain. Our ellipsis resolution strategy uses many pragmatic knowledge sources, including the information seeker's inferred task-related plan, his inferred beliefs, his anticipated discourse goals, and focusing strategies. However, we do not contend that a natural language system should use only pragmatic knowledge; a robust system will need to coordinate syntactic, semantic, and pragmatic techniques in order to fully understand the wide variety of elliptical utterances employed in human communication.

Real understanding necessitates identifying the speaker's discourse or communicative goal in uttering an elliptical fragment. We claim that a speaker who uses an elliptical fragment is relying heavily on discourse expectations and that such fragments are often used to accomplish discourse goals that must be marked if a complete sentence is used. Our work is the first to address the problem of recognizing the discourse goal fulfilled by an elliptical utterance. Our research shows that an ellipsis interpretation strategy controlled by discourse expectations and utilizing focusing heuristics facilitates recognition of the speaker's intent in uttering an intersentential elliptical fragment, including identification of his discourse goal.

ACKNOWLEDGMENTS

Some of this work was partially supported by a grant from the National Science Foundation, IST-8311400, and a subcontract from Bolt Beranek and Newman Inc. of a grant from the National Science Foundation, IST-8419162

I would like to thank Ralph Weischedel for his encouragement and suggestions and Lance Ramshaw for many helpful discussions. I also appreciate the comments and suggestions of Kathy McCoy, Alan Pope, Ralph Weischedel, and the anonymous reviewers on drafts of this paper.

REFERENCES

- Allen, James F. and Perrault, C. Raymond 1980 Analyzing Intention in Utterances. *Artificial Intelligence* 15: 143–178.
- Appelt, Douglas E. 1985 Planning English Referring Expressions. *Artificial Intelligence* 26(1):1–33.
- Austin, J. L. 1962 *How To Do Things With Words*. Harvard University Press, Cambridge, MA.
- Carberry, Sandra 1983 Tracking User Goals in an Information-Seeking Environment. In: *Proceedings of The National Conference on Artificial Intelligence*, Washington, DC: 59–63.
- Carberry, Sandra 1985 *Pragmatic Modeling in Information System Interfaces*. Ph.D. thesis, University of Delaware, Newark, DE.
- Carberry, Sandra 1988 Modeling the User's Plans and Goals. *Computational Linguistics Journal* 14(3):23–27.
- Carbonell, Jaime G. 1983 Discourse Pragmatics and Ellipsis Resolution in Task-Oriented Natural Language Interfaces. In: *Proceedings of the 21st Annual Meeting of the Association for Computational Linguistics*, Boston, MA: 164–168.
- Fikes, R. E. and Nilsson, N. J. 1971 STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving. *Artificial Intelligence* 2: 189–208.
- Finin, Tim and Drager, David 1986 GUMS: A General User Modeling System. In: *Proceedings of the Sixth Canadian Conference on Artificial Intelligence*, Montreal, Canada: 24–30.
- Flowers, M. and Dyer, M. E. 1984 Really Arguing with your Computer. *Proceedings of the National Computer Conference*; 653–659.
- Goodman, Bradley A. 1986 Reference Identification and Reference Identification Failures. *Computational Linguistics*; 273–305.
- Grice, H. Paul 1957 Meaning. *Philosophical Review* 56: 377–388.
- Grice, H. Paul 1969 Utterer's Meaning and Intentions. *Philosophical Review* 68: 147–177.
- Grice, H. Paul 1975 Logic and Conversation. In: Cole, P. and Morgan, J. L. (eds.), *Syntax and Semantics III: Speech Acts*, Academic Press, New York City, NY: 41–58.
- Grimes, J. E. 1975 *The Thread of Discourse*. Mouton.
- Grosz, Barbara J. 1977 The Representation and Use of Focus in a System for Understanding Dialogs. In: *Proceedings of the International Joint Conference on Artificial Intelligence*, Cambridge, MA: 67–76.
- Grosz, Barbara J. 1979 Utterance and Objective: Issues in Natural Language Processing. In: *Proceedings of the International Joint Conference on Artificial Intelligence*, Tokyo, Japan: 1067–1076.
- Grosz, Barbara and Sidner, Candace 1986 Attention, Intention, and the Structure of Discourse. *Computational Linguistics*, 12(3): 175–204.
- Hendrix, Gary G.; Scerdoti, Earl D.; Sagalowicz, Daniel; and Slocum, Jonathan 1978 Developing a Natural Language Interface to Complex Data. *ACM Transactions on Database Systems* 3(2): 105–147.
- Hobbs, Jerry R. 1979 Coherence and Coreference. *Cognitive Science* 3: 67–90.
- Kass, Robert and Finin, Tim 1987 Rules for the Implicit Acquisition of Knowledge About the User. In: *Proceedings of the Sixth National Conference on Artificial Intelligence*, Seattle, WA: 295–300.
- Labov, W. and Fanshel, D. 1977 *Therapeutic Discourse: Psychotherapy as Conversation*. Academic Press, New York City, NY.
- Litman, Diane J. 1986. Understanding Plan Ellipsis. In: *Proceedings of the Fifth National Conference on Artificial Intelligence*, Philadelphia, PA: 619–624.
- Litman, Diane and Allen, James 1987 A Plan Recognition Model for Subdialogues in Conversation. *Cognitive Science* 11:163–200.
- Mann, William; Morre, James; and Levin, James 1977 A Comprehension Model for Human Dialog. In: *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, Cambridge, MA: 77–87.
- McKeown, Kathleen R. 1985 *Text Generation*. Cambridge University Press, Cambridge, England.
- Perrault, R. and Allen, J. 1980 A Plan-Based Analysis of Indirect Speech Acts. *American Journal of Computational Linguistics*, 167–182.
- Pollack, M.; Hirschberg, J.; and Webber, B. 1982 User Participation in the Reasoning Processes of Expert Systems. In: *Proceedings of the Second National Conference on Artificial Intelligence*, (AAAL) Pittsburgh, PA: 358–361.
- Ramshaw, Lance 1989 A Metaplan Model for Problem-Solving Discourse. In: *Proceedings of the Fourth Conference of the European Chapter of the Association for Computational Linguistics*, Manchester, England.
- Reichman, Rachel 1978 Conversational Coherency. *Cognitive Science* 2:283–327.
- Reichman, Rachel 1984 Extended Person-Machine Interface. *Artificial Intelligence* 22:157–218.
- Robinson, Ann 1981 Determining Verb Phrase Referents in Dialogs. *American Journal of Computational Linguistics*; 1–18.
- Sacks, H.; Schegloff, E. A.; and Jefferson, G. 1974 A Simplest Systematics for the Organization of Turntaking for Conversation. *Language* 50: 696–735.
- Sidner, Candace L. 1981 Focusing for Interpretation of Pronouns. *American Journal of Computational Linguistics*; 217–231.
- Sidner, Candace L. 1983 What the Speaker Means: The Recognition of Speakers' Plans in Discourse] *Computers and Mathematics With Applications* 9(1):71–82.
- Sidner, Candace L. 1985 Plan Parsing for Intended Response Recognition in Discourse. *Computational Intelligence*; 1–10.
- Thompson, Bozena H. 1980 Linguistic Analysis of Natural Language Communication with Computers. In: *Proceedings of the 8th International Conference on Computational Linguistics*, Tokyo, Japan: 190–201.
- Waltz, David. 1978. An English Language Question Answering System for a Large Relational Data base. *Communications of the Association for Computing Machinery* 21(7):526–539.
- Weischedel, Ralph and Sondheimer, Norman K. 1982 An Improved Heuristic for Ellipsis Processing. In: *Proceedings of the 20th Annual Meeting of the Association for Computational Linguistics*, Toronto, Ontario, Canada: 85–88.

NOTES

1. Taken from Flowers and Dyer (1984).
2. A speech act is an action performed with an utterance, such as the acts of informing and making a promise, although other requirements may need to be met for the action to be successfully executed (Austin 1962).
3. Thanks to Ralph Weischedel for this example.
4. These transcripts were provided by the Department of Computer Science of the University of Pennsylvania.
5. These rules capture the behavior exhibited in the dialogs we studied. However, we have no experiments to indicate whether the model is psychologically valid.
6. Note that $\neg \text{Believe}(\text{IS}, \neg P) \rightarrow \text{Believe}(\text{IS}, P)$
7. This is similar to Grosz's focus spaces and the notion of an object being in implicit focus.
8. Some stack processing rules that were not discussed in Section 7.2 would influence the discourse stack but not affect how the elliptical fragments in the examples are understood. They have been omitted to keep the presentation short and to focus attention on the important aspects of the ellipsis processing strategy.