A PRAGMATICS-BASED APPROACH TO UNDERSTANDING INTERSENTENTIAL ELLIPSIS

Sandra Carberry
Department of Computer and Information Science
University of Delaware
Newark, Delaware 19715, USA

ABSTRACT

Intersentential elliptical utterances occur frequently in information-seeking dialogues. This paper presents a pragmatics-based framework for interpreting such utterances, including identification of the speaker's discourse goal in employing the fragment. We claim that the advantage of this approach is its reliance upon pragmatic information, including discourse content and conversational goals, rather than upon precise representations of the preceding utterance alone.

INTRODUCTION

The fragmentary utterances that are common in communication between humans also occur in man-machine communication. Humans persist in using abbreviated statements and queries, even in the presence of explicit and repeated instructions to adhere to syntactically and semantically complete sentences (Carbonell, 1983). Thus a robust natural language interface must handle ellipsis.

We have studied one class of elliptical utterances, intersentential fragments, in the context of an information-seeking dialogue. As noted by Allen(1980), such utterances differ from other forms of ellipsis in that interpretation often depends more heavily upon the speaker's inferred underlying task-related plan than upon preceding syntactic forms. For example, the following elliptical fragment can only be interpreted within the context of the speaker's goal as communicated in the first utterance:

[EX1] "I want to cash this check.
Small bills only."

Furthermore, intersentential fragments are often employed to communicate discourse goals, such as expressing doubt, which a syntactically complete form of the same utterance may not convey as effectively. In the following alternative responses to the initial statement by SPEAKER-1, F1 expresses doubt regarding the proposition stated by SPEAKER-1 whereas F2 merely asks about the jet's contents.

* This work has been 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

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

F1: "With 269 people on board?""

F2: "With infrared cameras on board?"

Previous research on ellipsis has neglected to address the speaker's discourse goals in employing the fragment but real understanding requires that these be identified (Mann, Moore, and Levin, 1977) (Webber, Pollack, and Hirschberg, 1982).

In this paper, we investigate a framework for interpreting intersentential ellipsis that occurs in task-oriented dialogues. This framework includes:

- [1] a context mechanism (Carberry, 1983) that builds the information-seeker's underlying plan as the dialogue progresses and differentiates between local and global contexts.
- [2] a discourse component that controls the interpretation of ellipsis based upon discourse goal expectations gleaned from the dialogue; this component "understands" ellipsis by identifying the discourse goal which the speaker is pursuing by employing the elliptical fragment, and by determining how the fragment should be interpreted relative to that goal.
- [3] an analysis component that suggests possible associations of an elliptical fragment with aspects of the inferred plan for the information-seeker.
- [4] an evaluation component which, given multiple possible associations of an elliptical fragment with aspects of the information-seeker's underlying plan, selects that association most appropriate to the discourse context and believed to be intended by the speaker.

INTERPRETATION OF INTERSENTENTIAL ELLIPSIS

As illustrated by [EX1], intersentential elliptical fragments cannot be fully understood in and of themselves. Therefore a strategy for interpreting such fragments must rely on knowledge obtained from sources other than the fragment itself. Three possibilities exist: the syntactic

** Taken from Flowers and Dyer(1984)

form of preceding utterances, the semantic representation of preceding utterances, and expectations gleaned from understanding the preceding discourse.

The first two strategies are exemplified by the work of Carbonell and Hayes(1983), Hendrix, Sacerdoti, and Slocum(1976), Waltz(1978), and Weischedel and Sondheimer(1982). Several limitations exist in these approaches, including an inability to handle utterances that rely upon an assumed communication of the underlying task and difficulty in resolving ambiguity among multiple interpretations. Consider the following two dialogue sequences:

SPEAKER: "I want to take a bus.
The cost?"

SPEAKER: "I want to purchase a bus.
The cost?"

If a semantic strategy is employed, the case frame representation for "bus" may have a "cost of bus" and a "cost of bus ticket" slot; ambiguity arises regarding to which slot the elliptical fragment "The cost?" refers. Although one might suggest extensions for handling this fragment, a semantic strategy alone does not provide an adequate framework for interpreting intersentential ellipsis.

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

Allen(1980) was the first to relate ellipsis processing to the domain-dependent plan underlying a speaker's utterance. Allen views the speaker's utterance as part of a plan which the speaker has constructed and is executing to accomplish his overall task-related goals. To interpret elliptical fragments, Allen first constructs a set of possible surface speech act representations for the elliptical fragment, limited by syntactic clues appearing within the fragment. The task-related goals which the speaker might pursue form a set of expectations, and Allen attempts to infer the speaker's goal-related plan which resulted in execution of the observed utterance. A part of this inference process involves determining which of the partially constructed plans connecting expectations (goals) and observed utterance are reasonable given the knowledge and mutual beliefs of the speaker and hearer. Allen selects the surface speech act which produced the most reasonable inferred plan as the correct interpretation.

Allen notes that the speaker's fragment must identify the subgoals which the speaker is pursuing, but claims that in very restricted domains, identifying the speaker's overall goal from the utterance is sufficient to identify the appropriate response in terms of the obstacles present in such a plan. For his restricted domain involving train arrivals and departures, Allen's interpreta-

tion strategy works well. In more complex domains, it is necessary to identify the particular aspect of the speaker's overall task-related plan addressed by the elliptical fragment in order to interpret it properly. More recently, Litman and Allen(1984) have extended Allen's model to a hierarchy of task-plans and meta-plans. Litman is currently studying the interpretation of elliptical fragments within this enhanced framework.

In addition to the syntactic, semantic, and plan-based strategies, a few other heuristics have been utilized. Carbonell(1983) uses discourse expectation rules that suggest a set of expected user utterances and relate elliptical fragments to these expected patterns. For example, if the system asks the user whether a particular value should be used as the filler of a slot in a case frame, the system then expects 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.

Real understanding consists not only of recognizing the particular surface-request or surface-inform, but also of inferring what the speaker wants to accomplish and the relationship of each utterance to this task. Interpretation of ellipsis based upon the speaker's inferred underlying task-related plan and discourse goals facilitates a richer interpretation of elliptical utterances.

REQUISITE KNOWLEDGE

A speaker can felicitously employ intersentential ellipsis only if he believes his utterance will be properly understood. 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 dialogue 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 discourse goals of the speaker; We claim that the requisite processing knowledge includes plan recognition strategies and focusing techniques.

1. Task-Related Plan

In a cooperative information-seeking dialogue, the information-provider is expected to infer the information-seeker's underlying task-related plan as the dialogue progresses. At any point in the dialogue, IS (the information-seeker) believes that some subset of this plan has been communicated to IP (the information-provider); therefore IS feels justified in formulating utterances under the assumption that IP will use this inferred task model to interpret utterances, including elliptical fragments.

An example will illustrate the importance of IS's inferred task-related plan in interpreting ellipsis. In the following, IS is considering purchase of a home mentioned earlier in the dialogue:

IS: "What elementary school do children in Rolling Hills attend?"

IP: "They attend Castle Elementary."

IS: "Any nearby swim clubs?"

An informal poll indicates that most people interpret the last utterance as a request for swim clubs near the property under consideration in Rolling Hills and that the reason for such an interpretation is their inference that IS is investigating recreational facilities that might be used if IS were to purchase the home. However, if we substitute the fragment

"Any nearby day-care centers?"

for the last utterance in the dialogue, then interpretation depends upon whether one believes IS wants his/her children to be bused, or perhaps even walk, to day-care directly from school.

2. Shared Beliefs

Shared beliefs of facts, 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 dialogue participants in a university domain know that each course has a teacher, or beliefs derived from the dialogue 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, he does believe it is mutually believed that IP believes this.

Understanding utterances 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 dialogue. The following example illustrates how IP's beliefs about IS influence understanding.

IS: "Who is teaching CS400?"

IP: "Dr. Brown is teaching CS400."

IS: "At night?"

The fragmentary utterance "At night?" is a request to know whether CS400 is meeting at night. However, if one precedes the above utterances with a query whose response informs IS that CS400 meets only at night, then the last utterance,

"At night?"

becomes an objection and request for corroboration or explanation. 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 already knows IP's beliefs regarding when CS400 meets, so a request for that information is not felicitous and a different intention or discourse goal is attributed to IS.

Allen and Perrault(1980) used mutual beliefs in their work on indirect speech acts and suggested their use in clarification and correction dialogues. Sidner(1983) models user beliefs about system capabilities in her work on recognizing speaker intention in utterances.

3. Anticipated Discourse Goals

The speaker's anticipated discourse goals form a third component of factual knowledge required for processing elliptical fragments. The dialogue preceding an elliptical utterance may suggest discourse goals for the speaker; these suggested discourse goals become shared knowledge between speaker and hearer. 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 dialogue:

IP: "Have you taken CS105 or CS170?"

IS: "At the University of Delaware?"

IP: "No, anywhere."

IS: "Yes, at Penn State."

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 are interpreted with the expectation that IS will eventually address this goal. IS's first utterance is interpreted as pursuing a discourse goal of seeking clarification of the question posed by IP; IS's last utterance answers the initial query posed by IP. However discourse expectations do not persist forever with intervening utterances.

4. Processing Knowledge

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 and focusing techniques are necessary in order to identify that portion of the underlying plan to which a fragmentary utterance refers.

Focusing mechanisms have been employed by Grosz(1977) in identifying the referents of definite noun phrases, by Robinson(1981) in interpreting verb phrases, by Sidner(1981) in anaphora resolution, by Carberry(1983) in plan inference, and by McKeown(1982) in natural language generation.

FRAMEWORK FOR PROCESSING ELLIPSIS

If an utterance is parsed as a sentence fragment, ellipsis processing begins. A model of any preceding dialogue contains a context tree (Carberry, 1983) corresponding to IS's inferred underlying task-related plan, a space containing IS's anticipated discourse goals, and a belief model representing IS's inferred beliefs.

Our framework is a top-down strategy which uses the information-seeker's anticipated discourse goals to guide interpretation of the fragment and relate it to the underlying task-related plan. The discourse component first analyzes the top element of the discourse stack and suggests potential discourse goals which IS might be expected to pursue. The plan analysis component uses the context tree and the belief model to suggest possible associations of the elliptical fragment with aspects of IS's inferred task-related plan. If multiple associations are suggested, the evaluation component applies focusing strategies to select the interpretation believed intended by the speaker --- namely, that most appropriate to the current focus of attention in the dialogue. The discourse component then uses the results produced by the analysis component to determine if the fragment accomplishes the proposed discourse goal; if so, it interprets the fragment relevant to the identified discourse goal.

PLAN-ANALYSIS COMPONENT

1. Association of Fragments

The plan-analysis component is responsible for associating an elliptical fragment with a term or conjunction of propositions in IS's underlying task-related plan. The analysis component determines, based upon the current focus of attention, the particular aspect of the plan highlighted by IS's fragment and the discourse goal rules infer how IS intends the fragment to be interpreted. 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).

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 has a limited set of valid instantiations within the existing plan. A constant associates with a term only if IP'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, JAN15)

the elliptical fragment

"February 2?"

will associate with this proposition only if IP 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. Although this discarding of possible associations does not occur often in interpreting elliptical fragments, actual human dialogues indicate that it is a real phenomenon. (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(_ss:&SECTIONS, SMITH)

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?"

is a variable term

_tme:&MTG-TMES

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.

2. Results of Plan-Analysis Component

The plan-analysis component constructs a conjunction of propositions PLPREDS and/or a term PLTERM representing that aspect of the information-seeker's plan highlighted by the elliptical fragment; STERM and SPREDS are produced by substituting into PLTERM and PLPREDS the terms in IS's fragment for the terms with which they are associated in IS's plan.

Figure 1: A Portion of the Expanded Context Tree for EXAMPLE-1

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 dialogue as retaining the fixed media specification:

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

"disks with two ports?"

We have noted the same phenomenon in a student advisement domain.

Thus when an elliptical fragment associates with a portion of the task-related plan or an expansion of one of its actions, the context established by the preceding dialogue must be used to replace information deleted from this streamlined, fragmentary utterance. The set of ACTIVE nodes in the context model form a stack of plans, the top-most of which is the current focused plan; each of these plans is the expansion of an action appearing in the plan immediately beneath it in this stack. These ACTIVE nodes represent the established global context within which the fragmentary utterance occurs, and the propositions appearing along this path contain information missing from the sentence fragment but presumed understood by the speaker.

If the elliptical fragment is a proposition, the analysis component produces a conjunction of propositions SPREDS representing that aspect of the plan highlighted by IS's elliptical fragment.

If the elliptical fragment is a constant, term, or term with attached propositions, the analysis component produces a term STERM associated with the constant or term in the fragment as well as a conjunction of propositions SPREDS. SPREDS consists of all propositions along the paths from the root of the context tree to the nodes at which an element of the fragment is associated with a plan element, as well as all propositions appearing along the previous ACTIVE path. The former represent the new context derived from IS's fragmentary utterance whereas the latter retain the previously established global context.

3. Example

This example illustrates how the plan-analysis component determines that aspect of IS's plan highlighted by an elliptical fragment. It also shows how the established context is maintained in interpreting ellipsis.

IS: "Is CS360 offered in Fall 1985?"

IP: "Yes."

IS: "Do any sections meet on Monday?"

IP: "One section of CS360 meets on Monday at 4PM
and another section meets on Monday at 7PM."

IS: "The text?"

A portion of IS's inferred task-related plan prior to the elliptical fragment is shown in Figure 1. Nodes along the ACTIVE path are marked by asterisks.

The semantic representation of the fragment
 "The text?"
 will be the variable term
   _book:&TEXTS
 This term associates with the term
   _txt:&TEXTS
 appearing at the node for the action

Learn-Text(IS,txt:&TEXTS)
   such that
 Uses(_ss:&SECTIONS,_txt:&TEXTS)

The propositions along the active path are

Course-Offered(CS360,FALL85)
 Is-Section-Of(_ss:&SECTIONS,CS360)
 Is-Offered(_ss:&SECTIONS,FALL85)
 Is-Syllabus-Of(_ss:&SECTIONS,_syl:&SYLBI)
 Teaches(_fac:&FACULTY,_ss:&SECTIONS)
 Is-Mtg-Day(_ss:&SECTIONS,MONDAY)
 Is-Mtg-Time(_ss:&SECTIONS,_tme:&MTG-TMES)
 Is-Mtg-Plc(_ss:&SECTIONS,_plc:&MTG-PLCS)

These propositions maintain the established context that we are talking about the sections of CS360 that meet on Monday in the Fall of 1985. The path from the root of the context model to the node at which the elliptical fragment associates with a term in the plan produces the additional proposition

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

The analysis component returns the conjunction of these propositions along with STERM, in this case
   _book:&TEXTS

The semantics of this interpretation is that IS is drawing attention to the term STERM such that the conjunction of propositions SPREDS is satisfied --- namely, the textbook used in sections of CS360 that meet on Monday in the Fall of 1985.

EVALUATION COMPONENT

The analysis component proposes a set of potential associations of the elliptical fragment with elements of IS's underlying task-related plan. The evaluation component employs focusing strategies to select what it believes to be the interpretation intended by IS --- namely, that interpretation most relevant to the current focus of attention in the dialogue.

We employ the notion of focus domains in order to group finely grained actions and associated plans into more general related structures. A focus domain consists of a set of actions, one of which is an ancestor of all other actions in the focus domain and is called the root of the focus domain. If an action is a member of a focus domain and that action is not the root action of another focus domain, then all the actions contained in the plan associated with the first action are also members of the focus domain. (This is similar to Grosz's focus spaces and the notion of an object being in implicit focus.)

The use of focus domains allows the grouping together of those actions that appear to be at approximately the same level of implicit focus

when a plan is explicitly focused. For example, the actions of learning from a particular teacher, learning the material in a given text, and attending class will all reside at the same focus level within the expanded plan for earning credit in a course. The action of going to the cashier's office to pay one's tuition also appears within this expanded plan; however it will reside at a different focus level since it does not come to mind nearly so readily when one thinks about taking a course.

The following are two of seven focusing rules used to select the association deemed most relevant to the existing plan context.

- [F1] Within the current focus space, prefer associations which occur within the current focused plan.
- [F2] Within the current focus space and current focused plan, prefer associations within the actions to achieve the most recently considered action.

DISCOURSE GOALS

We have analyzed dialogues from several different domains and have identified eleven discourse goals which occur during information-seeking dialogues and which may be accomplished via elliptical fragments. Three exemplary discourse goals are

- [1] 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.
- [2] Obtain-Corroboration: IS expresses surprise regarding some proposition P and requests elaboration upon and justification of it.
- [3] Seek-Clarify-Question: IS requests information relevant to clarifying a question posed by IP.

ANTICIPATED DISCOURSE GOALS

When IS makes an utterance, he is attempting to accomplish a discourse goal; this discourse goal may in turn predict other subsequent discourse goals for IS. For example, if IS asks a question, one anticipates that IS may want to expand upon his question. Similarly, utterances made by IP suggest discourse goals for IS. These Anticipated Discourse Goals provide very strong expectations for IS and may often be accomplished implicitly as well as explicitly.

The discourse goals of the previous section also serve as anticipated discourse goals. Three additional anticipated discourse goals appear to play a major role in determining how elliptical fragments are interpreted. One such anticipated discourse goal is:

Accept-Question: IP has posed a question to IS; IS must now accept the question either explicitly, implicitly, or indicate that he does not as yet accept it.

Normally dialogue participants accept such questions implicitly by proceeding to answer the question or to seek information relevant to formulating an answer. However IS may refuse 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 to discourse goals such as seeking confirmation, seeking the identity of an entity, seeking clarification of the posed question, or expressing surprise at the question.

THE DISCOURSE STACK

The discourse stack contains anticipated discourse goals which IS is expected to pursue. Anticipated discourse goals are pushed onto or popped from the stack as a result of utterances made by IS and IP. We have identified a set of stack processing rules which hold for simple utterances. Three examples of such stack processing rules are:

[SP1]When IP asks a question of IS, Answer-Question and Accept-Question are pushed onto the discourse stack.

[SP2]When IS poses a question to IP, Expand-Question is pushed onto the discourse stack. Once IP begins answering the question, the stack is popped up to and including the Expand-Question discourse goal.

[SP3]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.

The motivation for these rules is the following. When IP asks a question of IS, IS is first expected to accept the question, either implicitly or explicitly, and then answer the question. Upon posing a question to IP, IS is expected to expand upon this question with subsequent utterances or wait until IP produces an answer to the question. Although the strongest expectations are that IS will pursue a goal suggested by the top element of the discourse stack, this anticipated discourse goal can be passed over, at which point it no longer suggests expectations for utterances.

DISCOURSE INTERPRETATION COMPONENT

The discourse component employs discourse expectation rules and discourse goal rules. The discourse expectation rules use the discourse stack to suggest possible discourse goals for IS and activate the associated discourse goal rules. These discourse goal rules use the plan-analysis component to help determine the best interpretation of the fragmentary utterance relevant to the

suggested discourse goal. If a discourse goal rule succeeds in producing an interpretation, then the discourse component identifies that discourse goal and its associated interpretation as its understanding of the utterance.

1. Discourse Expectation Rules

The top element of the discourse stack activates the discourse expectation rule with which it is associated; this rule in turn suggests discourse goals which the information-seeker's utterance may pursue and activates these discourse goal rules. The following is an example of a discourse expectation rule:

[DE1]If the top element of the discourse stack is Answer-Question, then

1. Apply discourse goal rule DG-Answer-Quest to determine if the elliptical fragment is being used to accomplish the discourse goal of answering the question.
2. If no interpretation is produced, apply rule DG-Suggest-Answer-Question to determine if the elliptical fragment is being used to accomplish the discourse goal of suggesting an answer to the question.
3. If no interpretation is produced, apply discourse goal rule DG-Obtain-Info to determine if the elliptical fragment is being used to accomplish the discourse goal of seeking information in order to construct an answer to the posed question.

Once IS understands the question posed to him, IP's strongest expectation is that IS will answer the question; therefore first preference is given to interpretations which accomplish this goal. If IS does not immediately answer the question, then we expect a cooperative dialogue participant to work towards answering the question. This entails gathering information about the underlying task-related plan in order to construct a response.

2. Discourse Goal Rules

Discourse goal rules determine if an elliptical fragment accomplishes the associated discourse goal and, if so, produce the appropriate interpretation of the fragment. These discourse goal rules use the plan-analysis component to help determine the best interpretation of the fragmentary utterance relevant to the suggested discourse goal. However these interpretations are not actual representations of surface speech acts; instead they generally indicate elements of the plan whose values the speaker is querying or specifying. In many respects, this provides a better "understanding" of the utterance since it describes what the speaker is trying to accomplish.

The following is an example of a rule associated with a discourse goal suggested by the stack entry Accept-Response; the latter is pushed onto the discourse stack when IP responds to a question posed by IS.

DG-Obtain-Corrob

The discourse component calls the plan-analysis component to associate the elliptical fragment with a term STERM or a conjunction of propositions SPREDS in IS's underlying task-related plan. If IP believes it is mutually believed that IS already knows IP's beliefs about the value of the term STERM or the truth of the propositions SPREDS, then identify the elliptical fragment as accomplishing the discourse goal of expressing surprise at the preceding response; in particular, IS is surprised at the known values of STERM or SPREDS in light of the new information provided by IP's preceding response and the known aspect queried by IS's fragment.

The following is one of several rules associated with the discourse goal Answer-Question.

DG-Answer-Quest-2

If the elliptical fragment terminates with a period, then the discourse component calls the plan-analysis component to associate the elliptical fragment with a conjunction of propositions SPREDS in IS's underlying task-related plan. If successful, interpret the elliptical fragment as answering "Yes", with the restriction that the propositions SPREDS be satisfied in the underlying plan.

IMPLEMENTATION AND EXAMPLES

This pragmatics-based framework for processing intersentential ellipsis has been implemented for a subset of discourse goals in a domain consisting of the courses, policies, and requirements for students at a university. The following are working examples from this implementation.

The ellipsis processor is presented with a semantic representation of IS's elliptical fragment; it "understands" intersentential elliptical utterances by identifying the discourse goal which IS is pursuing in employing the fragment and by producing a plan-based interpretation relevant to this discourse goal.

EXAMPLE-1

This example illustrates a simple request for information.

IS: "Is CS360 offered in Fall 1985?"

IP: "Yes."

IS: "Do any sections meet on Monday?"

IP: "One section of CS360 meets on Monday at 4PM and another section meets on Monday at 7PM."

IS: "The text?"

Immediately prior to IS's elliptical utterance, the discourse stack contains the entries

Accept-Response Obtain-Information

The discourse goal rules suggested by Accept-Response do not identify the fragment as accomplishing their associated discourse goals, so the top entry of the discourse stack is popped; this indicates that IS has implicitly accepted IP's response. The entry Obtain-Information on the discourse stack activates the rule DG-Obtain-Info. Plan-analysis is activated to associate the elliptical fragment with an aspect of IS's task-related plan. The construction of STERM and SPREDS for this example was described in detail in the plan analysis section and will not be repeated here. Since our belief model indicates that IS does not currently know the value of STERM such that SPREDS is satisfied, this rule identifies the elliptical fragment as seeking information in order to formulate a task-related plan; in particular, IS is requesting the value of STERM such that SPREDS is satisfied --- namely, the textbook used in sections of CS360 that meet on Monday in the Fall of 1985.

EXAMPLE-2

This example illustrates an utterance in which IS is surprised by IP's response and seeks elaboration and corroboration of it. (The construction of SPREDS by the plan analysis component will not be described since it is similar to EXAMPLE-1.)

IS: "I want to take CS620 in Fall 1985.
Who is teaching it?"

IP: "Dr. Smith is teaching CS620 in Fall 1985."

IS: "What time does CS620 meet?"

IP: "CS620 meets at 8AM."

IS: "With Dr. Smith?"

IS's elliptical fragment will associate with the term

Teaches(_fac:&FACULTY,_ss:&SECTIONS)
in IS's task-related plan. SPREDS will contain the propositions

```
Course-Offered(CS620,FALL85)
Is-Section-Of(_ss:&SECTIONS,CS620)
Is-Offered(_ss:&SECTIONS,FALL85)
Is-Syllabus-Of(_ss:&SECTIONS,_syl:&SYLBI)
Teaches(SMITH,_ss:&SECTIONS)
Is-Mtg-Day(_ss:&SECTIONS,_day:&MTG-DAYS)
Is-Mtg-Time(_ss:&SECTIONS,_tme:&MTG-TMES)
Is-Mtg-Plc(_ss:&SECTIONS,_plc:&MTG-PLCS)
```

Immediately prior to the occurrence of the elliptical fragment, the discourse stack contains the entries

Accept-Response Obtain-Information

Accept-Response, the top entry of the discourse stack, suggests the discourse goals of 1) seeking confirmation or 2) seeking corroboration of a component of the preceding response or 3) seeking elaboration and corroboration of some aspect of this

Figure 2. A Portion of the Expanded Context Tree for EXAMPLE-3

response. The discourse goal rules Seek-Confirm and Seek-Identify fail to identify their associated discourse goals as accomplished by the user's fragment.

Our belief model indicates that IS already knows that SPREDS is satisfied; therefore the discourse goal rule DG-Obtain-Corrob identifies the elliptical fragment as expressing surprise at and requesting corroboration of IP's response. In particular, IS is surprised that SPREDS is satisfied and this surprise is a result of

- [1] the new information presented in IP's preceding response, namely that 8AM is the value of the term

```

      _tme:&MTG-TMES
in the SPREDS proposition
      Is-Mtg-Time(_ss:&SECTIONS,_tme:&MTG-TMES)

```

- [2] the aspect of the plan queried by IS's elliptical fragment, namely the SPREDS proposition

```

      Teaches(SMITH,_ss:&SECTIONS)

```

EXAMPLE-3

The following is an example which our framework handles but which poses problems for other strategies.

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

The first two utterances establish a plan context of late-registering, within which the elliptical fragment requests the fees involved in doing so. (Late registration generally involves extra charges.)

Figure 2 presents a portion of IS's underlying task-related plan inferred from the utterances

preceding the elliptical fragment. The parenthesized numbers preceding actions indicate the action's focus domain. IS's fragment associates with the term

```

      _lreg:&MONEY

```

in IS's inferred plan, as well as with terms elsewhere in the plan. However none of the other terms appear in the same focus space as the most recently considered action, and therefore the association of the fragment with

```

      _lreg:&MONEY

```

is selected as most relevant to the current dialogue context. The discourse stack immediately prior to the elliptical fragment contains the single entry

Provide-For-Assimilation

This anticipated discourse goal suggests the discourse goals of 1)providing further information for assimilation and 2)seeking information in order to formulate the task-related plan. The utterance terminates in a "?", ruling out provide for assimilation. Therefore rule DG-Obtain-Info identifies the elliptical fragment as seeking information. In particular, the user is requesting the fee for late registration, namely, the value of the term

```

      _cst1:&MONEY

```

such that SPREDS is satisfied, where SPREDS is the conjunction of the propositions

```

Course-Offered(_crs:&COURSE,_sem:&SEMESTERS)
Is-Section-Of(_ss:&SECTIONS,_sem:&SEMESTERS)
Is-Offered(_ss:&SECTIONS,_sem:&SEMESTERS)
Costs(LATE-REG,_cst1:&MONEY)

```

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 partially due to the fact that our mechanisms for retaining dialogue context are based upon the view 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 dialogue context will fail to identify "teacher of CS263" as the information requested by IS.

One might argue that the elliptical fragment in the above dialogue relies heavily upon the syntactic representation of the preceding utterance and thus a syntactic strategy is required for interpretation. This may be true. However if we view dialogues such as the above as investigating task-related plans in a kind of "breadth-first" fashion, then 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 the plan-based framework can be extended to handle many such dialogues, perhaps by using meta-plans to represent how IS is constructing his task-related plan.

CONCLUSIONS

This paper has described a pragmatics-based approach to interpreting intersentential elliptical utterances during an information-seeking dialogue in a task domain. Our framework coordinates many knowledge sources, including the information-seeker's inferred task-related plan, his inferred beliefs, his anticipated discourse goals, and focusing strategies to produce a rich interpretation of ellipsis, including identification of the information-seeker's discourse goal. This framework can handle many examples which pose problems for other strategies. We claim that the advantage of this approach is its reliance upon pragmatic information, including discourse content and conversational goals, rather than upon precise representations of the preceding utterance alone.

ACKNOWLEDGEMENTS

I would like to thank Ralph Weischedel for his encouragement and direction in this research and Lance Ramshaw for many helpful discussions and suggestions.

REFERENCES

1. Allen, J.F. and Perrault, C.R., "Analyzing Intention in Utterances", *Artificial Intelligence*, 15(3), 1980
2. Carberry, S., "Tracking User Goals in an Information-Seeking Environment", AAAI, 1983
3. Carberry, S., "Pragmatic Modeling in Information System Interfaces", forthcoming Ph.D. Dissertation, Dept. of Computer Science, University of Delaware, Newark, Delaware
4. Carbonell, J.G., and Philip Hayes, "Recovery Strategies for Parsing Extragrammatical Language", *Amer. Journal of Comp. Ling.*, Vol.9, No.3-4, 1983
5. Carbonell, J.G., "Discourse Pragmatics and Ellipsis Resolution in Task-Oriented Natural Language Interfaces", *Proc. 21st Annual Meeting of ACL*, 1983
6. Flowers, M. and M.E. Dyer, "Really Arguing With Your Computer", *Proc. of Nat. Comp. Conf.*, 1984
7. Grice, H.P., "Meaning", *Phil. Rev.* 66, 1957
8. Grice, H.P., "Utterer's Meaning and Intentions", *Phil. Rev.*, 68, 1969
9. Grosz, B.J., "The Representation and Use of Focus in a System for Understanding Dialogs", *IJCAI*, 1977
10. Grosz, B.J., Joshi, A.K., and Weinstein, S., "Providing a Unified Account of Definite Noun Phrases in Discourse", *Proceedings 21st Annual Meeting of ACL*, 1983
11. Hendrix, G.G., E.D. Sacerdoti, and J. Slocum, "Developing a Natural Language Interface to Complex Data", *SRI International*, 1976
12. Litman, D.J., and Allen, J.F., "A Plan Recognition Model for Clarification Subdialogues", *Proceedings of the International Conference on Computational Linguistics*, 1984
13. Mann, W., J. Moore, and J. Levin, "A Comprehension Model for Human Dialogue", *IJCAI*, 1977
14. McKeown, K.R., "The Text System for Natural Language Generation: An Overview", *Proc. of the 20th Annual Meeting of ACL*, 1982
15. Perrault, C.R., and Allen, J.F., "A Plan-Based Analysis of Indirect Speech Acts", *American Journal of Computational Linguistics*, July 1980
16. Robinson, A.E., "Determining Verb Phrase Referents in Dialogs", *American Journal of Computational Linguistics*, Jan. 1981
17. Sidner, C.L., "What the Speaker Means: The Recognition of Speakers' Plans in Discourse", *Comp. and Maths. with Appls.*, Vol.9, No.1, 1983
18. Sidner, C.L., "Focussing for Interpretation of Pronouns", *American Journal of Computational Linguistics*, Oct. 1981
19. Waltz, D.L., "An English Language Question Answering System for a Large Relational Data Base", *Comm. of ACM*, vol21, No.7, 1978
20. Webber, B.L., M.E. Pollack, and J. Hirschberg, "User Participation in the Reasoning Processes of Expert Systems", *Proc. of Nat. Conf. on Art. Int.*, 1982
21. Weischedel, R.M. and N. Sondheimer, "An Improved Heuristic for Ellipsis Processing", *Proc. 20th Annual Meeting of ACL*, 1982