

USER MODELS: THE PROBLEM OF DISPARITY

Sandra Carberry

Department of Computer and Information Sciences
University of Delaware
Newark, Delaware 19716, U.S.A.

ABSTRACT

A significant component of a user model in an information-seeking dialogue is the task-related plan motivating the information-seeker's queries. A number of researchers have modeled the plan inference process and used these models to design more robust natural language interfaces. However in each case, it has been assumed that the system's context model and the plan under construction by the information-seeker are never at variance. This paper addresses the problem of disparate plans. It presents a four phase approach and argues that handling disparate plans requires an enriched context model. This model must permit the addition of components suggested by the information-seeker but not fully supported by the system's domain knowledge, and must differentiate among its components according to the kind of support accorded each component as a correct part of the information-seeker's overall plan. It is shown how a component's support should affect the system's hypothesis about the source of error once plan disparity is suggested.

1. INTRODUCTION

Communication as we know it involves more than simply answering isolated queries. When two individuals participate in an information-seeking dialogue, the information-provider uses the context within which each query occurs to interpret the query, determine the desired information, and formulate an appropriate response. This context consists of more than mere knowledge of the previous questions and answers. A cooperative participant uses the information exchanged during the dialogue and knowledge of the domain to hypothesize a model of the speaker; this model is adjusted and expanded as the dialogue progresses and is called a *user model*.

Perhaps the most significant component of a user model is the listener's belief about the underlying task motivating the information-seeker's queries and his partially developed plan for accomplishing this task. A number of researchers have modeled the plan inference process [Allen 1980], [Carberry 1983], [Grosz 1977], [Litman 1984], [Perrault 1980], [Robinson 1981], [Sidner 1983], and these models have been used to understand indirect speech acts [Perrault 1980], provide helpful responses [Allen 1980], interpret pragmatically ill-formed queries [Carberry 1986], understand intersentential ellipsis [Allen 1980, Carberry 1985], and identify the kind of response intended by a speaker [Sidner 1983].

However in each case, four critical assumptions have been made:

- [1] The information-seeker's knowledge about the task domain may be lacking but is not erroneous.
- [2] The information-seeker's queries never address aspects of the task outside the system's knowledge. Such systems maintain the closed world assumption [Reiter 1978].

- [3] The information provided by the information-seeker is correct and not misleading.
- [4] The underlying plan inferred by the system prior to analysis of a new utterance is a partially instantiated version of the plan under consideration by the information-seeker.

These assumptions eliminate the possibility that the information-seeker might ask queries irrelevant to the task at hand, that the information-seeker might ask about details outside the system's limited knowledge, that the information-seeker might accidentally provide misleading information, and that the system might have made erroneous inferences from previous queries. The end result is that the system believes that the underlying task-related plan inferred by the system and the task-related plan under construction by the information-seeker are never at variance with one another.

If we want systems capable of understanding and appropriately responding to naturally occurring dialogue, natural language interfaces must be able to deal with situations where those assumptions are not true. Our analysis of transcripts of naturally occurring information-seeking dialogues indicates that human participants attempt to detect inconsistencies in the models and repair them whenever possible. We claim that natural language systems must do likewise; otherwise they will be unable to respond appropriately and cooperatively to dialogue that humans regard as natural.

This paper presents a taxonomy of disparate plan models, according to how the model inferred by the information-provider reflects the information-seeker's model of his task. We claim that plan inference must be extended to include a four phase approach to handling disparate plans and that this approach requires a richer model than maintained by current systems. We show how the support that an information-provider accords a component as a correct part of the model affects her hypothesis about the source of error once plan disparity is suggested.

2. TYPES OF MODELS

An information-seeking dialogue contains two participants, one seeking information and the other attempting to provide that information. Underlying such a dialogue is a task which the information-seeker wants to perform, generally at some time in the future. The information-seeker poses queries in order to obtain the information necessary to construct a plan for accomplishing this task. Examples of such tasks include pursuing a program of study in a university domain, treating a patient in a medical domain, and taking a vacation in a travel domain.

A cooperative natural language system must attempt to infer the underlying task-related plan motivating the information-seeker's queries and use this plan to provide cooperative, helpful responses [Carberry 1983, 1985]. We call the system's model of this plan a *context model*. A context model is one component of a user model.

We are concerned here with cases in which the system's context model fails to mirror the plan under construction by the information-seeker. Disparate plan models may be classified according to how the model inferred by the system differs from the information-seeker's model of his task:

- [1] *erroneous models*, representing cases in which the model inferred by the system is inconsistent with the information-seeker's model. If the information-seeker were to examine the system's model in such cases, he would regard it as containing errors.
- [2] *overly-specialized models*, representing cases in which the model inferred by the system is more restricted than that intended by the information-seeker.
- [3] *overly-generalized models*, representing cases in which the model inferred by the system is less specific than that intended by the information-seeker.
- [4] *knowledge-limited models*, representing cases in which the model inferred by the system fails to mirror the plan under construction by the information-seeker, due to the system's limited domain knowledge.

The use of default inferencing rules may produce erroneous or overly-specialized models. Erroneous models may also result if the information-seeker's statements are inaccurate or misleading or if the system uses focusing heuristics to relate new utterances to the existing plan context. Overly-generalized models may result if the information-seeker fails to adequately communicate his intentions (or the system fails to recognize these intentions). Knowledge-limited models may result if the information-seeker's domain knowledge exceeds that of the system.

A fifth category, *partial models*, represents cases in which the system has inferred only part of the information-seeker's plan; subsequent dialogue will enable the system to further expand and refine this context model as more of the information-seeker's intentions are communicated. We do not regard partial models as disparate structures: were the information-seeker to examine the system's inferred partial plan, he would regard it as correctly modeling his intentions as communicated in the dialogue thus far.

3. RELATED WORK

Several research efforts have addressed problems related to plan disparity. Kaplan[1982] and McCoy[1986] investigated misconceptions about domain knowledge and proposed responses intended to remove the misconception. However such misconceptions may not be exhibited when they first influence the information-seeker's plan construction; in such cases, disparate plans may result and correction will entail both a response correcting the misconception and further processing to bring the system's context model and the plan under construction by the information-seeker back into alignment.

Pollack[1986] is studying removal of what she terms the "appropriate query assumption" of previous planning systems; she proposes a richer model of planning that explicitly reasons about the information-seeker's possible beliefs and intentions. Her overall goal is to develop a better model of plan inference. She addresses the problem of queries that indicate the information-seeker's plan is inappropriate to his overall goal, and attempts to isolate the erroneous beliefs that led to the inappropriate query. This is a subclass of "erroneous plans", since upon hearing the query, the system should detect that its context model no longer agrees with that of the information-seeker. However, queries deemed inappropriate by the system may signal phenomena other than inappropriate user plans. For example, the information-seeker may have shifted focus to another aspect of the overall task without successfully conveying this to the system, the

information-seeker may be addressing aspects of the task outside the system's limited knowledge, or the system's context model may have been in error prior to the query.

Pollack is concerned with issues that arise when the information-seeker's plan is incorrect due to a misconception. She assumes that, immediately prior to the user making the "problematic" query, the system's partial model of the user's plan is correct. We argue that since the system's inference mechanisms are not infallible and communication itself is imperfect, the system must contend with the possibility that its inferred model does not accurately reflect the user's plan. Previous research has failed to address this problem.

4. PROBLEM POSED BY DISPARATE MODELS

Grosz[1981] claimed that communication can proceed smoothly only if both dialogue participants are focused on the same subset of knowledge. Extending this to inferred plans, we claim that communication is most successful when the information-provider's and information-seeker's models mirror one another. But clearly it is unrealistic to expect that these models will never diverge, given the different knowledge bases of the two participants and the imperfections of communication via dialogue. Thus the information-provider (IP) and the information-seeker (IS) must be able to detect inconsistencies in the models whenever possible and repair them. Clearly a natural language system must do the same.

This view is supported by the work of Pollack, Hirschberg, and Webber[1982]. They conducted a study of naturally occurring expert-novice dialogues and suggested that such interaction could be viewed as a negotiation process, during which not only an acceptable solution is negotiated but also understanding of the terminology and the beliefs of the participants. The context model is one component of IP's beliefs, as is her belief that it accurately reflects the plan under construction by IS.

5. AN APPROACH TO DISPARATE MODELS

A study of transcripts of naturally occurring information-seeking dialogues indicates that humans often employ a four phase approach in detecting and recovering from disparate plan structures. Therefore a natural language interface that pursues the same strategy will be viewed as acting naturally by human users. The next sections discuss each of these phases.

5.1. DETECTION AND HYPOTHESIS FORMATION

As claimed earlier, since IP is presumed to be a cooperative dialogue participant, IP must be on the lookout for plan disparity. We have identified three sources of clues to the existence of such disparity:

- [1] the discourse goals of IS, such as expressing surprise or confusion
- [2] relevance of IS's current utterance to IP's inferred model
- [3] focus of attention in the model

IS can express surprise or confusion about IP's response, thereby cuing the possibility of plan disparity. Consider for example the dialogue presented in Figure 1. This dialogue was transcribed from a radio talk show on investments* and will be referred to as the "IRA example"; utterances are numbered for later reference. Plan disparity is suggested when IS, in utterance [5], expresses confusion at IP's previous response.


On the other hand, IS's query may contradict or appear irrelevant to what IP believes is IS's overall task, leading IP to suspect that her context model may not reflect IS's plan. Or IS's

* Transcripts of these dialogues were provided by the Department of Computer Science of the University of Pennsylvania

- [1] IS: "I'm a retired government employee but I'm still working. I'd like to start out an IRA for myself and my wife --- she doesn't work."
- [2] IP: "Did you work outside of the government last year?"
- [3] IS: "Yes I did."
- [4] IP: "There's no reason why you shouldn't have an IRA for last year."
- [5] IS: "I thought they just started this year."
- [6] IP: "Oh no. IRA's were available as long as you are not a participant in an existing pension."
- [7] IS: "Well, I do work for a company that has a pension."
- [8] IP: "Ahh. Then you're not eligible for 81."

Figure 1. Individual Retirement Account Dialogue^{3*}

query may require so sharp an unsignaled shift in focus as to cause IP to be suspicious; the strongest expectations are for speakers to address aspects of the task closely related to the current focus of attention [Sidner 1981, McKeown 1985, Carberry 1983]. The dialogue presented in Figure 2, and henceforth referred to as the "Kermit example", illustrates a case in which plan disparity is suggested by an abrupt shift in focus of attention. Upon completion of utterance [4], IP's model of IS's plan might be represented as



Both humans and machines have limited knowledge. Suppose that IP does not know how to purchase floppy disks. Then from IP's limited knowledge, IS's next query,

"How late is the University Bookstore open?"

will not appear to address an aspect of the plan inferred for IS, or any expansion of it. IP could just respond by

- [1] answering the direct question, if possible, ignoring its ramifications
- [2] responding "I don't know", if the direct answer is not available

However cooperative human information-providers are expected to try to understand the import of a query and provide as cooperative a response as they can.

Grice's maxim of relation [Grice 1975] suggests that IS believes the query to be relevant to the overall dialogue. Several possibilities exist. IS may be shifting focus to some aspect of a higher-level task that includes transferring files as a subaction. One such higher-level task might be to compose a document using the SCRIBE text formatting system, and the aspect queried by the new utterance might be the purchase of a SCRIBE manual from the university bookstore; in this case, the subtask of the overall task represented by the existing context model might be

* Minor alterations have been made to the dialogue to remove restarts and extraneous phrasing.

- [1] IS: "I wish I could transfer files between the Vax and my PC."
- [2] IP: "Kermit lets you do that."
- [3] IS: "How do I get Kermit?"
- [4] IP: "The computing center will give you a copy if you bring them a floppy disk."
- [5] IS: "How late is the University Bookstore open?"

Figure 2. File Transfer Via Kermit Dialogue

the transfer of files containing the document so that they can be modified using a PC editor.

On the other hand, focusing heuristics and the absence of discourse markers [Sidner 1985] suggest that the new query is most likely to be relevant to the current focus of attention. So IP should begin trying to determine how IS's utterance might relate to the currently focused subtask in the context model, and consider the possibility that IS's domain knowledge might exceed IP's or might be erroneous.

5.2. RESPONSE PHASE

Webber[1986] distinguishes between answers and responses. She defines an answer as the production of the information or execution of the action requested by the speaker but a response as

"the respondent's complete informative and performative reaction to the question which can include ... additional information provided or actions performed that are salient to this substitute for an answer."

Our analysis of naturally occurring dialogue indicates that humans *respond*, rather than *answer*, once disparate models are detected. These responses often entail additional actions, including a negotiation dialogue to ascertain the cause of the discrepancy and enable the models to be modified so that they are once again in alignment. A robust natural language interface must do the same, since the system must have an accurate model of the information-seeker's plan in order for cooperative behavior to resume.

The appropriate response depends on the cause of the discrepancies. In the case of a knowledge-limited model, IP should attempt to understand IS's utterance in terms of IP's limited knowledge and provide any pertinent helpful information, but inform IS of these limitations in order to avoid misleading IS by appearing to implicitly support his task-related plan.

Consider again our example of file transfer via Kermit, presented in Figure 2. We assume that, in addition to a domain-dependent set of plans, IP's knowledge base contains a generalization hierarchy of actions and entities.

Suppose that IP's knowledge base contains the plans

- ```

To: Have(<agent>:PERSON, <x>:BOOK)
 Action: Purchase(<agent>, <x>)

To: Purchase(<agent>:PERSON, <x>:TEXTBOOK)
 Action: GoTo(<agent>, <p>:BOOKSTORE, <t>:TIME)
 where Sells(<p>, <x>)
 Between(<t>, <t1>:TIME, <t2>:TIME)
 Opens(<p>, <t1>)
 Closes(<p>, <t2>)

```

IP can reason that IS's last query is relevant to a plan for purchasing a textbook at the bookstore. This is simple plan inference

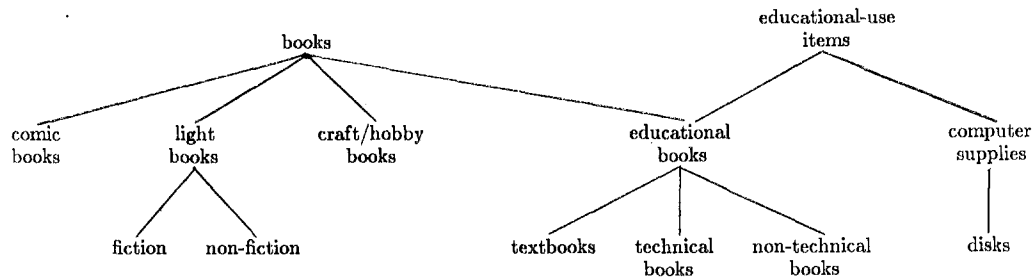


Figure 3. Object Taxonomy for Kermit Example

as embodied in our TRACK system [Carberry 1983]. However IP cannot connect purchasing a book with her model of IS. So IP may begin trying to expand on her knowledge. Suppose that IP's taxonomy of objects is as shown in Figure 3 and that IP's domain knowledge includes the existence of many instances of <u>, <v>, <w>, and <x> such that

```

Sells(UDEL-BOOKSTORE, <u>:NOVEL)
Sells(UDEL-BOOKSTORE, <v>:TECHBOOK)
Sells(UDEL-BOOKSTORE, <w>:NONTECHBOOK)
Sells(UDEL-BOOKSTORE, <x>:TEXTBOOK)

```

Novels are a subclass of light-books and, technical-books, non-technical-books, and textbooks are subclasses of educational-books. But educational-books are a subclass of educational-use-items, as are floppy disks. Thus IP can generalize textbooks to educational-use-items, note that this class also contains disks, and then hypothesize that perhaps IS thinks that the bookstore sells floppy disks, since it sells other educational-use items. This reasoning might be represented by the rule

```

If Class-1 is a subclass of Class-2, and for many of the
other subclasses of Class-2 there exist many members
<y> such that
 P(...,<y>),
then one can hypothesize that perhaps there exists
<z> such that
 P(...,<z>:Class-1)

```

This rule can be applied in the absence of contradictory domain knowledge. Having thus hypothesized that perhaps

```
Sells(UDEL-BOOKSTORE, <z>:DISK)
```

from

```

Sells(UDEL-BOOKSTORE, <v>:TECHBOOK)
Sells(UDEL-BOOKSTORE, <w>:NONTECHBOOK)
Sells(UDEL-BOOKSTORE, <x>:TEXTBOOK)

```

IP can hypothesize the higher-level goals

```
Purchase(IS, <z>:DISK)
```

```
Have(IS, <z>:DISK)
```

the last of which is a component of IP's model of IS.

Since IP has constructed a plan that may reasonably be ascribed to IS, is relevant to the current focus of attention, and about which IP's knowledge is neutral, IP can hypothesize that the cause of the plan disparity may be that IS has more extensive domain knowledge. IP can now respond to IS. This reply should of course contain a direct answer to IS's posited question. But

this alone is insufficient. In a cooperative information-seeking dialogue, IS expects IP to assimilate the dialogue and relate utterances to IS's inferred underlying task in order to provide the most helpful information. If IP limits herself to a direct response, IS may infer that IP has related IS's current utterance to this task and that IP's knowledge supports it --- that is, that IP also believes IS can purchase a floppy disk at the bookstore. Joshi's revised maxim of quality [Joshi 1983] asserts that IP's response must block false inferences. In addition, as a helpful participant, IP should include whatever evidence IP has for or against the plan component proposed by IS. An appropriate response would be:

"The University Bookstore is open until 4:30 PM. But I don't know whether it sells floppy disks. However it does sell many other items of an educational nature, so it is perhaps a good place to try."

The above example concerned a knowledge-limited model caused by IP's limited domain knowledge. Other kinds of models suggest different reasoning and response strategies. If IP has failed to make the inferences IS assumed would be made, then subsequent utterances by IS may appear appropriate to a more specific model than IP's current model. Earlier, we referred to this class as overly-generalized models. In these cases, IP may enter a clarification dialogue to ascertain what IS intends.

In other cases, such as when overly-specialized or erroneous models are detected, a negotiation dialogue must be initiated to "square away" [Joshi 1983] the models; otherwise, IS will lack confidence in the responses provided by IP (and therefore should not continue the dialogue), and IP will lack confidence in her ability to provide useful replies (and therefore cannot continue as a cooperative participant). As with any negotiation, this is a two-way process:

- [1] IP may select portions of the context model that she feels are suspect and justify them, in an attempt to convince IS that IS's plan needs adjustment, not IP's inferred model of that plan.
- [2] IP may formulate queries to IS in order to ascertain why the task models diverge and where IP's model might be in error.

The IRA example illustrates a negotiation dialogue. In utterance [6], IP selects a suspect component of her context model and provides justification for it. IS's next utterance informs IP that the assumption on which this component was based is incorrect; IP then notifies IS that IP recognizes the error and that her context model has been repaired. The information-seeking dialogue then resumes.

### 5.3. MODEL RECONSTRUCTION

Once the cause of model disparity is identified, IP and IS must adjust their models to remove the disparities. Once again,

this depends on the cause of the disagreement. In the case of a knowledge-limited model, IP should incorporate the components she believes to be part of IS's plan structure into her context model, noting however that her own knowledge offers only limited support for them. In this way, IP's model reflects IS's, enables IP to understand (within her limited knowledge) how IS plans to accomplish his objectives, and permits IP to use this knowledge to understand subsequent utterances and provide helpful information.

If IP's model is in error, she must alter her context model, as determined through the negotiation dialogue. She may also communicate to IS the changes that she is making, so that IS can assure himself that the models now agree. On the other hand, if IS's model is in error, IP may inform IS of any information necessary for IS to construct an appropriate plan and achieve his goals.

#### 5.4. SUMMARY

The arguments in the preceding sections are based on an analysis of transcripts of human information-seeking dialogues and indicate that an appropriate approach for handling the plan disparity problem entails four phases:

- [1] detection of disparate models
- [2] hypothesis formation as to the cause of the disparities
- [3] extended response, often including a negotiation dialogue to identify the cause of the disparities
- [4] model modification, to "square away" the plan structures.

Since this approach is representative of that employed by human dialogue participants, a natural language interface that pursues the same strategy will be viewed as acting naturally by its human users.

#### 6. ENRICHED CONTEXT MODEL

The knowledge acquired from the dialogue and how it was used to construct the context model are important factors in detecting, responding to, and recovering from disparate models. Human dialogue participants typically employ various techniques such as focusing strategies and default rules for understanding and relating dialogue, but they appear to have greater confidence in some parts of the resultant model than others. Natural language systems must employ similar mechanisms in order to do the kind of inferencing expected by humans and provide the most helpful responses. We claim that the representation of the inferred plan must differentiate among its components according to the support which the system accords each component as a correct and intended part of the inferred plan. This view parallels Doyle's Truth Maintenance System [Doyle 1979], in which attitudes are associated with reasons justifying them.

We see four kinds of support for plan components:

- [1] whether the system has inferred the component directly from what IS said.
- [2] whether the system has inferred the component on the basis of its own domain knowledge, which the system cannot be certain IS is aware of.
- [3] the kinds of mechanisms used to add each component to the model, (for example, default rules that select one component from among several possibilities, or heuristics that suggest a shift in focus of attention), and the evidence for applying the mechanism.
- [4] whether the system's domain knowledge supports, contradicts, or is neutral regarding inclusion of the component as part of a correct overall plan.

The first three are important factors in formulating a hypothesis regarding the source of disparity between the system's model and IS's plan. If the system believes that IS intends the

system to recognize from IS's utterance that G is a component of IS's plan, then the system can add G to its context model and have the greatest faith that it really is a component of IS's plan. Therefore such components are unlikely sources of disparity between the system's context model and IS's plan.

Components that the system adds to the context model on the basis of its domain knowledge will be strongly believed by the system to be part of IS's plan, but not as much as if IS had directly communicated them. These components resemble "keyhole recognition" rather than "intended recognition" [Sidner 1985, 1983]. Since IS may not have intended to communicate them, they are more likely sources of error than components which IS intended IP to recognize.

Consider for example a student advisement system. If only BA degrees have a foreign language requirement, the query

"What course must I take to satisfy the foreign language requirement in French?"

may lead the system to infer that IS is pursuing a Bachelor of Arts degree. If only BS degrees require a senior project, then a subsequent query such as

"How many credits of senior project are required?"

suggests plan disparity. Either the second query is inappropriate to IS's overall goal [Pollack 1986] or the system's context model is already in error. Since the component

Obtain-Degree(IS, BACHELOR-OF-ARTS)

was inferred on the basis of the system's domain knowledge rather than directly from IS's utterance, it is suspect as the source of error.

The mechanisms used to add a component to the context model affect IP's faith in that component as part of IS's overall plan. Consider again the IRA example in Figure 1. In utterance [4], IP has applied the default assumption that IS was not covered by a pension program during the year in question (at that time, rules on IRAs were different). IS's next utterance expresses confusion at IP's response, thereby cuing the possibility of plan disparity. In utterance [6], IP selects the component added to the context model via the default assumption as a possible source of the disparity, tells IS that it is part of IP's context model, and attempts to justify its inclusion.

Analysis of naturally occurring dialogues such as that in Figure 1 indicate that humans use mechanisms such as default inference rules and focusing heuristics to expand the context model and provide a more detailed and tractable arena in which to understand and respond to subsequent utterances. Natural language systems must use similar mechanisms in order to cooperatively and naturally engage in dialogue with humans. However these rules select from among multiple possibilities and therefore produce components that are more likely sources of error than components added as a result of IS's direct statements or inferences made on the basis of the system's domain knowledge.

The fourth kind of differentiation among components --- whether the system's domain knowledge supports, contradicts, or is neutral regarding inclusion of the component as part of a correct overall plan --- is important in recovering from disparate plans. Even an expert system has limited domain knowledge. Furthermore, in a rapidly changing world, knowledgeable users may have more accurate information about some aspects of the domain than does the system. For example, a student advisement system may not be altered immediately upon changing the teacher of a course. Thus we believe that the context model must allow for inclusion of components suggested by the information-seeker, including whether the system's domain knowledge contradicts, supports, or is neutral regarding the component.

For example, upon determining that IS's domain knowledge may exceed the system's in the Kermit dialogue, the system should expand its existing model to incorporate the acquired knowledge about how IS believes floppy disks can be obtained. The plan components creatively constructed can be added to the system's model, but as components proposed by IS and not fully supported by the system's knowledge. In this manner, the system can assimilate new utterances that exceed or contradict its limited domain knowledge and develop an expanded context model which serves as "knowledge" that can be referred back to in the ensuing dialogue.

## 7. SUMMARY

This paper has addressed the problem of disparity between the context model inferred by a natural language system and the plan under construction by an information-seeker. We have presented a four phase approach and have argued that handling disparate plans requires an enriched context model. This model must permit the addition of components suggested by the information-seeker but not fully supported by the system's domain knowledge and must differentiate among its components according to the kind of support accorded each component as a correct part of the information-seeker's overall plan. We have further argued that support for a component should affect the system's hypothesis about the source of error once plan disparity is suggested.

## 8. ACKNOWLEDGEMENTS

I want to thank Joe Brady, Kathy Cebulka, Dan Chester, Kathy McCoy, Martha Pollack, and Ralph Weischedel for their many helpful discussions and comments on this work, and Dan Chester and Kathy McCoy for their comments and suggestions on this paper.

## 9. REFERENCES

- Allen, James F., "Analyzing Intention in Utterances", *Artificial Intelligence* 15(3), 1980
- Carberry, Sandra, "Pragmatic Modeling in Information System Interfaces", Ph.D. Dissertation, Department of Computer Science, University of Delaware, 1985
- Carberry, Sandra, "Tracking User Goals in an Information-Seeking Environment", *Proceedings of the National Conference on Artificial Intelligence*, 1983
- Carberry, Sandra, "Using Inferred Knowledge to Understand Pragmatically Ill-Formed Queries", to appear in *Communication Failure in Dialogue*, Ronan Reilly editor, North Holland, 1986
- Doyle, Jon, "A Truth Maintenance System", *Artificial Intelligence* 12(3), 1979
- Grice, H. P., "Logic and Conversation", In *Syntax and Semantics*, Cole and Morgan, editors, Academic Press, 1975
- Grice, H. P., "Utterer's Meaning and Intentions", *Philosophical Review* 68, 1969
- Grice, H. P., "Meaning", *Philosophical Review* 56, 1957
- Grosz, Barbara, "Focusing and Description in Natural Language Dialogues", in *Elements of Discourse Understanding*, Joshi, A., Webber, B., and Sag, I., editors, Cambridge University Press, 1981
- Grosz, Barbara, "The Representation and Use of Focus in a System for Understanding Dialogs", *Proceedings of the International Joint Conference on Artificial Intelligence*, 1977
- Joshi, Aravind K., "Mutual Beliefs in Question-Answer Systems", *Mutual Knowledge*, Academic Press, 1983
- Kaplan, S. Jerrold, "Cooperative Responses from a Portable Natural Language Query System", *Artificial Intelligence* 19, 1982
- Litman, Diane J. and Allen, James F., "A Plan Recognition Model for Clarification Subdialogues", *Proceedings of the International Conference on Computational Linguistics*, 1984
- McCoy, Kathleen F., "Generating Responses to Property Misconceptions Using Perspective", to appear in *Communication Failure in Dialogue*, Ronan Reilly, Editor, 1986
- McKeown, Kathleen R., *Text Generation*, Cambridge University Press, 1985
- Perrault, C. R. and Allen, J. F., "A Plan-Based Analysis of Indirect Speech Acts", *American Journal of Computational Linguistics*, 1980
- Pollack, Martha, "Inferring Domain Plans in Question-Answering", forthcoming Ph.D. Dissertation, University of Pennsylvania, 1986
- Pollack, Martha, "Some Requirements for a Model of the Plan-Inference Process in Conversation", to appear in *Communication Failure in Dialogue*, Ronan Reilly, editor, North Holland, 1986
- Pollack, Martha, Hirschberg, Julia, and Webber, Bonnie, "User Participation in the Reasoning Processes of Expert Systems", *Proceedings of the National Conference on Artificial Intelligence*, 1982
- Reiter, Ray, "On Closed World Data Bases", *Logic and Data Bases*, Gallaire, H. and Minker, J., editors, Plenum Press, 1978
- Robinson, Ann E., "Determining Verb Phrase Referents in Dialogs", *American Journal of Computational Linguistics*, 1981
- Sidner, Candace L., "Plan Parsing for Intended Response Recognition in Discourse", *Computational Intelligence* 1(1), 1985
- Sidner, Candace L., "What the Speaker Means: The Recognition of Speakers' Plans in Discourse", *Computers and Mathematics With Applications* 9(1), 1983
- Sidner, Candace L., "Focusing for Interpretation of Pronouns", *American Journal of Computational Linguistics*, 1981
- Webber, Bonnie L., "Questions, Answers, and Responses: Interacting with Knowledge Base Systems", to appear in *On Knowledge Base Management Systems*, M. Brodie and J. Mylopoulos, editors, Springer-Verlag, 1986