RECOGNIZING INTENDED MEANING AND SPEAKERS' PLANS


Candace L. Sidner and David J. Israel

Bolt Beranek and Newman Inc.
10 Houlton St.
Cambridge, MA 02138



ABSTRACT

Human conversational participants depend upon the ability of their partners to recognize their intentions, so that those partners may respond appropriately. In such interactional the speaker can encode his intentions that the hearer act in a variety of sentence types. Instead of telling the hearer what to do. the speaker may Just state his goals, and expect a response that meets these goals. This paper presents a new model for recognizing the speaker's Intended meaning in determining a response. We show that this recognition makes use of the speaker's plan, his beliefs about the domain and about the hearer's relevant capacities.*


## 1. STATEMENT OP THE PROBLEM

Human conversational participants depend upon the ability of their partners to recognize their intentions, so that those partners may be capable of responding appropriately. For instance, in the dialogue below,

D1-1 SI: I want to see the drawing of the new design layout.
2 *S2:* OK. Here it is. (shows sheet with new design)
3 S1: There isn't room to put in the color code charts at the bottom of the picture. Can you move up the main layout?
4 S2: Sure, I'll bring back the new design in half an hour.

the speaker's desires are encoded in a variety of sentence types. Instead of telling the hearer what to do, the speaker states his goals, and expects a response that meets them, at least partway. This paper presents a model in which recognizing the speaker's intended meaning plays a fundamental part in determining a response.


*Acknowledgements: The research reported In this paper was supported in part by the Advanced Researoh Projects Agency and was monitored by the Office of Naval Research under Contract No. N00014-77-C-0378.


## 2. INTRODUCTION

One of the goals of the Knowledge Representation and Natural Language group at BBN has been to provide powerful general tools for natural language processing and to build a language understanding system for a decision maker using a graphics display. We envision a decision maker aocessing information from a database that can be represented visually: he/she needs to collect information from the database, add to it, change it and define new features.

A special feature of this language understanding system is the assumption that the human user expresses him/herself naturally. He/she can utter more than the direct imperatives and can ask questions besides the direct questions in typical of most current AI language systems.

Natural communication with a partner is possible, however, only if the latter can reason both about the intent behind the speaker's utterances and about Its own responses to those utterances. Such reasoning Involves - at least - (1) bringing to bear the kinds of knowledge people have before they enter into a given discussion, and (2) making use of the knowledge they gain in the discussion.

An organizing framework for a system that can reason about the the speaker's Intentions has been explored in our group and is reported elsewhere [Brachman et al. 1979]. In that framework we experimented with an Implementation of Allen's [1979] model of a plan-based approach to speech act recognition. This model provided us with a perspective within which to use models of speaker's beliefs and wants [Cohen, 1978]) as well as a framework for reasoning about the speaker's plans. We found that we needed to expand on Allen's model in two ways: (1) by recognizing a richer form of plans, and (2) by making explicit the connection between the speaker's Intentions as structured by his plans and the response intended by the speaker.

In this paper we will describe the new theoretical model for recognizing speaker's intended meaning and the computational tools we are currently using to implement that model. We think the model is particularly powerful because it is the basis of a system that oan reason about the kinds of planning bugs the acknowledgement of, and the response to whioh are often refleoted in conversational exchanges.

## 3. DEPININO INTENDED SPEAKER MEANING

Our goal is to provide a computational model of the hearer's interpretation of the speaker's intended meaning. The <u>intended meaning</u> of an utterance we define as that set of <propositional attitude (e.g., belief, want, intend), propoaitional content> pairs that the speaker wants to Induce in the hearer by means of the utterance.

The notion of the intended meaning of an utterance oan be Illustrated by contrasting it with that of semantic meaning. The semantic meaning of a declarative utterance is the propositlonal content assigned to the type of the utteranoe by the semantic rules of the language. For instance, if someone says, "You're a prince," the semantic meaning is (roughly) that the person addressed by the speaker la tne son of a king. By contrast, the intended meaning is dependent on (inter alia) the individual psychological state of the speaker at the time and place of utterance. The speaker may mean that he thinks the hearer is a really nice guy and wants to tell him so, or he may be saying something quite different. The speaker, using irony, may mean that the addressee is just the opposite of a nice guy.

This (in some ways extreme) example demonstrates that speaker's intended meaning, though correlated, is not in general identical, with semantic meaning. Comprehending the semantic meaning of the utterance forms the oasis for discerning the intended meaning, but the latter (typically) also requires the use of the following kinds of belief on the part of the hearer:

1. Beliefs about the characteristics of the current situation,

2. Beliefs about the speaker's beliefs and goals,

3. Beliefs about the context of discussion (the discourse context) as a special aspect of (1),

4. Beliefs about what conventions for action exist between the speaker and hearer,*

5. Beliefs about what is mutually believed, as between speaker and hearer, with respect to (1) through (4).

A sample exchange will indicate the role of these kinds of belief. In the example below, the user is interacting with our system to display some information (in this case, a portion of an ATN network). The user's first two utterances are simple direct imperatives that indicate that the user wants the system to display a part of the net and then move the focus to a subpart of the display.

D2-1 U: Display the clause level network.
2 S: <display of network> OK.
3 U: Now focus on the pre verbal constituents.
4 S: <display of subnet> OK.
5 U: No, I want to see S/AUX.

What does the user mean by her third utterance (utteranoe 5)? The answer depends on what she believes about the net objects to which she has referred. Suppose she thinks that S/AUX is part of the preverbal constituents. Then she is communicating that the display la wrong and what's wrong with it; she Intends for S/AUX to be Included in the display with the other constituents. Suppose, alternatively that she thinks that S/AUX is not part of the preverbal constituents. She is still indicating that she wants to see S/AUX, but also that she has changed her mind about the display in some way and intends S/AUX to be visible. [As to the nature of her change of mind, this depends on whether she has realized the error of her ways, or at least the mismatch between her views on S/AUX and the system's.]

In ohoosing among its available responses, the system must utilize its model of the user's beliefs about the domain and its model of what the user takes to be mutually believed between the two of them about that domain. For example, the user might have thought that S/AUX was one of the preverbal constituents, and thought the system belleved this also. She would then have expected and Intended the system to include that state in the display. If the user had been right about this, the system would indeed have included it. But the user's "No" indicates to the system some bug in her plan, a bug stemming either from a faulty model of the domain itself or from faulty expectations about the system's model. (For simplicity, we assume that the system is omniscient about the ATN grammar.]

If, on the other hand, the system doesn't oonclude that the user takes S/AUX to be among the preverbal constituents, and if it believes that she takes that to be mutually believed, then the system must again use its models of her and of her model of itself to determine what action is intended by the user [E.g., should it compress the current display to make room for S/AUX; should it erase the current display and bring up a new one, centered on S/AUX, etc.?]. This decision may depend on the kinds of conventions alluded to in (4) above. In general, of course. people's behavior in conversational situations also depends on the relative status of the conversational partners, on what the participants think will benefit themselves, as well as not harm others, and the like. These social considerations are significant to human interaction, but for the remainder of this paper, we'll assume that the system responds in a slavishly cooperative way, that la, it has no interest beyond serving the user.

There are two ways to view the intentions of another agent. The first is simply in terms of one's beliefs about what the other person wants and believes. This is keyhole recognition (see Cohen et al [198l]). One person decides what he thinks another intends simply by observing him through a keyhole. (E.g. I decide that you are looking for your umbrella, on the basis of your looking around the room wltn your coat on, when I believe you believe that it's raining outside). Keyhole recognition of a user's wants is central to Genesereth's MACSYMA advisor [1978]; it also forms the basis of plan recognition in both Schmidt et al'a BELIEVER [19781 system and in Wilensky's story understanding work L1978].

The intended recognition of what someone is doing, on the other hand, la relevant for communicative situations [Orlce 1957, Allen.1979]. A speaker says something to a hearer, and Intends that the hearer recognize the intention that lies behind the utterance. The speaker is attempting to "give the hearer a piece of its mind" and It's essential to the success of the speaker's attempt that the hearer recognize it as such.

## 4. A MODEL OF RECOGNITION OF INTENDED MEANING

The hearer's taak in recognizing what the apeaker meant by an utteranoe la to be understood as follows:

1. to produce an explanation for the utterance, stated in terms of the speaker's beliefs and wants.

2. to use the explanation aa a basis for a response.

We use the term "explanation" because the hearer Is trying to answer the question "Why did the speaker say that to me?* The answer to this question - the proffered explanation of the speaker's act in uttering what he/she did - in turn produces new beliefs about the speaker; theae will form part of the basis of the hearer's response.

The explanation, in general, will have the form of a set of pairs of propositlonal attitudes and propositlonal contents attributed by the hearer to the speaker. [E.g. <bellef, that S7AUX is part of

•This, of course, is relevant only to a special class of situations; a olass which Includes the kind of interaction the BBN system must handle.

the preverbal oonstituents> <want, that I display all ooaponents of the preverbal conaituents>. etc.j Certain beliefs play a central role in explaining why the speaker said what he did:

Explanatory beliefs

1. beliefs about the speaker's goal and the plan to achieve it,

2. beliefs about the hearer's capacities,

3. beliefs about the hearer's dispositions to act given information about the speaker's wants.

The problem we pose for ourselves is determining how to infer beliefs of these kinds, and how to use them to distinguish, between Intended, and helpful, but unintended, responses. We want our system to recognize and produce the intended response whenever possible, and to be able to produce a helpful response when appropriate.

To model the construction of the required explanation, we begin with Grice's theory of speaker meaning [1957, 1969]. Grice notes that there are certain kinds of evidence which are normally available to an audience on the basis of which the audience (is intended) to draw certain conclusions about the speaker's intended meaning. These include the features of the utterance, mappings between those features and proposltional attltude-proposltlonal content pairs that the audience, assumed to be a competent speaker/hearer of the language, is supposed to be able to, and is intended to grasp, etc. For example, the feature: DECLARATIVE will be mapped to the speaker's wanting the hearer to believe the speaker believes the proposltional content of the utterance; while imperatives will be mapped to the speaker's wanting the hearer to believe the speaker wants the hearer to bring about the state of affairs expressed by the proposltional component of the utterance.

Somewhat more formally: an audience, for the utterance of a certain sentence S1, who la believed by the speaker to have certain attributes A, is expected to be able to recognize certain features of the utterance and to be able to draw from those features certain conclusions about what the speaker intended in uttering S1 in that context. [One such audience attribute, of course. Is competence in the language of S1: others are both more interesting and more situation-specific.] These conclusions are (or at least inolude):

Intended Conclusions

1. S1 has certain features (call them F1 ... Fn).

2. S1 is correlated, in virtue of such features and the rules of the language, with the pair (p , PC(SD).

3. The speaker intends the audience to believe that the speaker p's that PC(S1).

4. By sincerity (see below), the speaker does p that PC(S1),

5. the speaker intended that the hearer p* that PC(S1).

[In the above, "p" is a schematic letter which takes verbs of proposltional attitude as substituends; "PC", a schematic letter which takes declarative sentences as substituends.] We oan apply this theory directly to the sample dialogues. For example, let us consider a sample utterance from the dialogue D2, understood, however, as the initial utterance of a discourse:

SI: I want to see S/AUX.

Intuitively, we would like the theory to allow us to show how an audience (even a computer system) would conclude that the user wants to see S/AUX, and that the user wants it to believe that he/she has this desire.

The set of relevant features F, attributes A and mappings C are (or anyway include):

- o F1 = S1 is in declarative mood
- o F2 = S1 was uttered intentionally by U
- o F3 = S1 was intentionally directed at S
- o A1 = s is a computer system with a graphics display, and U knows this
- o A2 = S believes U Is sincere
- o C1 s F1 maps to U's wanting the intended audience to believe that U believes that U wants to see S/AUX.

Our proposed system will make default assumptions guaranteeing F2, F3, A1 and A2, recognize that F1, and apply cl to S1. The system can then utilize the intended conclusions and infer directly that:

1. U Intended (S to reoognlze) that S1 is correlated with U's wanting S to believe that U believes that U wants to see S/AUX (derived from Intended conclusion 3 and C1).

2. By sincerity: U believes that U wants to see S/AUX.

3. By reliability: U wants to see S/AUX.*).

This, of course, is what, on intuitive grounds, wanted the system to conclude.

•Actually the hearer may be intended to have a different proposltional attitude p' toward a related proposition. For simplicity, we'll assume these are the same.

•Simply stated these rules, for the case of belief, are: Sincerity! If x wants y to believe that x believes that q, then x believes that q. Reliability: If y believes that x believes that q and that x is reliably informed about q, then y will believe that q. The basis for these rules is the intuition that the speaker is sincere about his beliefs, and that what he believes he believes reliably, at least for certain subjeot matters, such as his own present state of mind.

## 5. EXTENDING THE MODEL

While the Gricean framework provides a starting point for recognizing the speaker's communicative intentions, it, of course, does not provide a recipe for inferring the intended response. Given, e.g., that the user wants the system to believe that the user wants to see S/AUX, and nothing more, the system could simply say "Yes, I understand," (or "Gee, let me add that to my data base of beliefs about you") - a behavior the user probably did not intend. At the same time, the system could decide to provide a lot of information by showing the whole ATN network and highlighting S/AUX. Such behavior might even be helpful; but it is not, we can presume, the intended response.

To determine the response the user intended, the system must consider the utterance in a larger situational context. This context is determined by what (it thinks) the user is doing, what (it thinks) the user thinks the system can do, and how cooperative (it thinks) the user takes the system to be. We now turn to a description of a method for inferring the intended response from the initial intended conclusions about the user's beliefs and desires.

We have augmented the Gricean framework to enable the system to derive a situation-specific explanation for the user's having the wants and beliefs s/he is believed to have. In particular, the system can be viewed as asking itself for an explanation of (some of) the beliefs it attributes to the user. The explanation is of the same type as that given earlier. For example, to explain why the user wants the system to perform some action, the system would infer that the user is pursuing a plan in which that action is a step. [This process must stop somewhere; in fact it stops because some plans are simply assumed to be entered into for their own sake and to require no further explanation.]

An example will illustrate what we have in mind. For utterance S1 above and the conclusions about the user's wants regarding S/AUX given previously, the system seeks to explain why the user wants the system to believe the user wants to see S/AUX, and perhaps why the user wants to see S/AUX. To answer the first, the system determines if any of the plans it has provisionally attributed to the user contains this step; and if so, it determines what relevant capacities the user believes the system to have. For the case at hand, since there are many such plans (deleting S/AUX, re-arranging its arcs, etc.) and since this is the initial interaction, no detailed plan-information is deducible. The user is assumed to believe, however, that the system has capacities relevant to seeing S/AUX--e.g., displaying it on the screen. The system concludes the user intends this capacity to be used, and since the system is cooperative, the system produces a display. An explanation for the user's wanting to see S/AUX may not be forthcoming. (It may also not be needed for the system's response planning.)

This extended theory depends not only on the Gricean framework, but also on the ability to create an explanation based on the user's plans. This last involves determining the user's goals by grasping the intended meaning of his utterance, (where the goals are structured in a hierarchy of goals and sub-goals), and using the speaker's recognized goals as an expectation model for the remaining part of the discourse.

To implement this model, we are using a number of available AI tools (the implementation is not complete). The system must have definitions of a number of plans, so we are using Sacerdoti-based procedural networks of plans [Sacerdoti, 1977]. Beliefs and wants must also be represented, and for this we are relying on Allen's and Cohen's models of belief and want contexts. A crucial aspect of this model is a method of "parsing" the user's wants as steps in plans; we are currently studying algorithms using an ATN formalism, but modified to allow for bugs in a plan, recognizable with a small bug library (see [Sussman, 1975]). This model bears some similarity to Genesereth's [1978] plan recognizer; it is distinguished by recognizing a different class of bugs than Genesereth's. This method makes it possible to use a plan, once selected from the collection of plans by unique sub-steps, as an expectation device for the remaining part of a discourse. Finally, we use standard antecedent reasoning for deducing the correlations between utterance features and propositional attitudes, and for relating user plans and the system's capacities.

## 6. REASONING ABOUT USER'S "BUGGY" PLANS

In the previous sections, we have shown the utility of explanations reflecting (among other things) beliefs about the speaker's goals and about his beliefs about the system's capacities. In this section we will demonstrate what additional reasoning such a model enables. In particular, we will show that such explanations provide a system with a means of discerning bugs in a user's plan. In the first example, the user is unaware of a bug in his plan. The system, after recognizing the bug, must inform the user, because no satisfactory response is possible until the bug is resolved. In the second example, the user discerns a bug and informs the system; awareness of this bug allows the system to recognize the intended meaning of a subsequent utterance. We will present enough detail of each example to permit the reader to see how a program could embody this reasoning process.

Let us return to example D2 given below.

D2-1 U: Display the clause level network.
2 S: <display of net> OK.
3 U: Now focus on the preverbal constituents.
4 S: <display of subnet consisting of S/Q, S/HOW, S/, S/QDET, S/NP, S/DCL.> OK.
5 U: No, I want to see S/AUX.

After the request to focus on the preverbal constituents, the system recognizes that the user's plan involves examining the preverbal constituents:*

Examine User "preverbal constituents"
Cause User
(Display System "preverbal constituents")
See User "preverbal constituents"

The plan has two steps; the first is to cause the system to perform a display and the second is to look at the displayed items. [Plans typically have preconditions; steps may be primitive, or may be composed of actions, requiring other plans as well.]

Following the last utterance of D2, the system can infer that by "no," the user is signifying that his plan has failed in some way. In interpreting the rest of the sentence, the system will reason about the user's intentions (as shown in the previous section) and conclude that the user's intention is to see S/AUX. This intention nearly matches step 2 of the plan deduced. It differs because S/AUX is not a part of the preverbal constituents. Using a small bug library, the system will recognize a possible bug in the Examine

*We assume here that the phrase "preverbal constituents" is interpreted appropriately, but will not discuss this interpretation here.

plan.

To account for the bug, the system can reason in either of two ways*. On the one hand, If it now has reason to believe that the user believes that the preverbal constituents Include S/AUX, It will conclude that there is a bug in the user's plan. This is a private (not mutual) belief; but it prevents the system from responding in the way intended by the user (to display S/AUX) for not enough of the user's intentions are clear to decide how to do the display (e.g., to include S/AUX or to show It alone). Henoe the system will respond by indicating what the bug la and by asking about the particular mode of display desired.

On the other hand, if the system believes that the user believes the preverbal constItutents are (mutually believed to be) disjoint from S/AUX, then the system will oonolude that the user has scrapped his current plan, and that this conclusion Is one the system Is Intended to deduce. In this case, displaying S/AUX Is the Intended response, but the system must still ask how to display It, since It Is not clear whether the user Intended It to be displayed alone or with the subnet. A person. In such circumstances, would probably conclude that S/AUX should be displayed alone, because s/he oould deduce that In general If a plan Is scrapped, effects of Its partial realization are no longer desired. However, this heuristic may be too general for systems which still have limited reasoning capacities, and hence we have chosen not to Include such rules.

This example demonstrates two aspects of our system: Its use of plan-attributions, Inferred In the course of Interpreting the user's Intentions, to recognize bugs In plans, and its use of private as well as mutual beliefs to determine its response when what the user intends Is unclear due to a "buggy" plan.

The next example Is a variation of D1. The task here is to Interact with a user when a graphics display is available for representing information about a database. The database Is represented in KL-OWE [Brachman et al, 1979], and the display consists of a graphic representation of XL-ONE concepts.

D3-1 U: I want to see the generic concept named employee.
2 *S:* OK. (displays concept In mid-screen)
3 U: I can't fit a new individual concept below it.
4 Can you move it up?
5 S: Sure. (moves up the generic concept)

The system's problem In responding to "Can you move it up?" la to determine whether the user meant his utterance directly as a question about the system's abilities, or intended rather to be taken to be directing the system to move the conoept under discussion. Its decision depends on inferring the speaker's plan and, in particular, on what it believes the user's model of its own capacities to be.

This example illustrates a feature of natural interchanges: a user may have a plan in mind, and carry out a part of it, without considering possible undesired side effect; when one occurs, it may be recognized and eliminated. In D3 the user is carrying out the plan of accessing the concept for employee so that a/he can add a new employee to the database. S/he wants the system to display the employee concept, but has not foreseen that its display looatlon might be infelicitous. After the inappropriateness is discovered, the user indloatea the difficulty and expects it to be corrected. Just how the bug in D3 is corrected depends on whether the user already believes the system can move things up and intends the system to do so, or has to find this out first.

From the system's point of view, the decision about what the user means may cause it to respond differently in various cases. Suppose the system thInjcs the user believes that the system can move up concepts on the screen. Then when the user indloatea that his plan has a flaw (D3-3), the system must conclude that the user's plan is blocked by the lack of space for a new concept. When the question about moving the employee concept is raised, the system will oonolude that the user intends to tell the system to perform the move by asking about a precondition of the action s/he wants, a precondition which consists in the system's having a capacity It is mutually believed to possess. The system is intended to move up the conoept, not simply to answer the question.

A different scenario is as follows. Suppose the system thinks the user is uaware that the system is oapable of moving up the conoept. Then when the user Indicates that his plan has a flaw and asks about moving the employee concept, the system will conclude that the user Intends to find out whether it has that ability, as part of finding a means of resolving the bloolc. In this case, if the system moves the concept, that is a bit of helpful behavior, one not Intended to be recognized as Intended by the user.

Ve will outline in some detail how our system reasons in such contexts by showing what plans are deduced, what rules are needed and how the reasoning proceeds in the case of D3-3 and 4. The relevant user plan is:

Add-Data <User> <netplece> <data> <acreen-looation>
Consider-aapect <User> <netpiece>
Put User <data> at <aoreen-location>

The Add-Data plan states that to add data, a user must consider some aspect of a network part (netplece) and then put some data at a screen location. Even after recognizing that the user wants some data displayed after D3-1, the system can not deduoe that add-data is the user's plan. Since there are many ways to oonsider some aspect of a net (ask for a display, think about it, ask to be Informed about its contents), as well as many other plans for which displaying a netpiece is a first step, the user cannot Be understood to have intended the system to recognize that his plan was to Add-Data . All the system can oonolude is that the user wants the employee concept displayed, and it responds accordingly.

In reasoning about D3-3, "I can't fit a new individual oonoept below it," the system concludes that among the speaker's intentions mutually presumed to be recognized is that the user produoed a declarative utteranoe with the propositional content that the user cannot fit a new individual conoept (abbrievate e2) below the generic concept (abbreviated e1):

BELIEF1; (Say User System (Declarative
(Not (Can User (Pit User e2
(below e1))))))

Prom this, the system ooncludes that the user wants the system to believe that the user believes that it can't fit e2 below e1, and that the user in fact believes that he can't. The system then infers the embedded proposition [(Not (Can ...))], and that the user intended that that proposition be mutually believed. Using a (default type) rule to the effect that whenever a user says that it can't

•Actually there is a third case—the system was wrong about what "preverbal constituents" refers to. As mentioned above, we Ignore this case.

bring about a certain atata-of-affaira orparforma cartain action, tha User la tailing tha ayataa that It wanta that stata-of-affairs brought about, tha ayataa concludaa that It la intandad to baliava that tha uaar wanta it to baliava

WANT1:(Fit User e2 (below e1))

The system seeks a partial explanation of this intention. It decides that the previous request for a display of the generic concept, together with the inferred intention to fit, are good evidence for the add-data plan as the intended plan.

Now the system can bring to bear a rule to the effect that if it believes that the user has informed it that he/she can't perform a certain action which he/she wants to perform as part of some plan he/she is pursuing, then it should conclude that the user intends that that action should be unblocked. Whether and how the system is expected to respond depends upon whether the system believes that the user believes there is some action available to the system relevant to this unblocking. In fact, the system might have several relevant capacities (such as moving up screen objects or erasing the screen as alternative ways to make room). Where the system believes the user knows this and hence concludes that the user wants to exploit some system capacity or other, it must await further information to determine which action was meant. In any case, on the basis of attributing the unblock plan, it can interpret the user's question as a way of bringing about a move-up action rather than simply a desire for information. Even if no such request as "Can you move it up?" were to follow, the system would have a basis for asking about the user's intent ("Do you want me to move up the concept or empty the screen?").

If the system believed the user was unaware of its capacities to move screen objects up, it would reason no further on D3-3 (again, because it has not recognized any intention on the part of the speaker that it act). D3-4 allows the system to deduce that the user wants to know if the system can move e1 up, since yes-no questions are taken as signalling intentions to knowif. The system, in seeking an explanation, can conclude that this is the first step in finding an agent with a certain capacity, and that this action is a means of unblocking the step of putting objects on the screen. The system must respond to the user's intention by telling the user it can move up the display, but the choice actually to move is made as helpful behavior* because no intention that the system move anything up has been recognized.

In both of the above cases, the plan to Add-Data to the screen is known to be in effect. Hence once the bug is cleared away, the system is prepared to interpret subsequent utterances in light of this plan. Were the user to say**: "Now put an ic with a name role with value 'Jon' on the screen," the system would recognize this request as part of Add-Data and determine the intended location for the ic on the screen--below the generic concept for employee. In this way, the plan becomes an expectation device for the next portion of the conversation.

This example illustrates not only that mutual ballafa about tha ayataa'a oapaoltlaa affaot tha ayataa*a dataralnatlon of tha uaar'a intantiona, but alao that tha full axplanation of aaon uttaranoa daapana tha ayataa*a undaratandlng of tha uaar*a goala and aubaaquant uttaranoaa.

•baaad on rulaa about ha1ping out whan ona haa an approprlata capacity. Aa with othar hauriatica, it ▪ay ba wiaa to aonltor oarafully what tha ayataa doas to ba halpful alnoa aoaa halpful actiona, if not In faot daalrad by tha apaakar can ba aaally undooa, but otnara hava non-trivial aida~affacta.

••Ic» and "rola with valua" ara KL-OWE taraa

## 7. CONCLUSIONS

In this paper we have attempted to demonstrate the centrality of the problem of determining a speaker's intended meaning in the context of designing a natural language interface for a graphics display facility. We have shown the several different kinds of evidence necessary for such a determination: beliefs about the current situation, the speaker's goals, the discourse context, the speaker's model of the hearer's capacities, and others. We have sketched both a theoretical model and a computational framework for realizing that model. The task of the system is to infer the user's intended meaning and to use the explanations it generates in doing so to determine the response intended by the user. Our system rests on many available AI tools, and extends yet others.

## 8. REFERENCES


Allen, J. [1979]. "A Plan Based Approach to Speech Act Recognition." University of Toronto, Department of Computer Science TR No. 131.

Brachman, R. [1978]. "A Structural Paradigm for Representing Knowledge." BBN Report No. 3605, Cambridge, MA.

Brachman, R. [1979], R. Bobrow, P. Cohen, J. Klovstad, B. Webber, W. Woods. "Research in Natural Language Understanding: Annual Report." BBN Report No. 4747, Cambridge, MA.

Cohen, P. [1978]. "On Knowing What to Say: Planning Speech Acts." University of Toronto, Department of Computer Science TR No. 118.

Cohen, P., R. Perrault and J. Allen. "Beyond Question Answering." to appear in Lehnart and Ringle (eds.), Strategies for Natural Language Processing, Hillsdale, NJ: Lawrence Erlbaum Assoc., 1981.

Genesereth, M. [1978]. "Automated Consultation for Complex Computer Systems." PhD. Disseration, Harvard University, Cambridge, MA.

Grice, H.P. [1957]. "Meaning." Philosophical Review, 66(3), p. 377-388.

Grice, H.P. [1969]. "Utterer's Meaning and Intentions." Philosophical Review, 68(2), p. 147-177.

Sacerdoti, E. [1977]. "A Structure for Plans and Behavior." American Elsevier, NY.

Schmidt, C.F. [1978], N.S. Sridharan and J.L.Goodson. "The Plan Recognition Problem: An Intersection of Psychology and Artificial Intelligence." Artificial Intelligence, 11(1-2), p. 45-83.

Sussman, G.J. [1975]. "A Computer Model of Skill Acquisition." American Elsevier, N.Y.

Wilensky, R. [1978]. "Understanding Goal Based Stories." Yale University Research Report No. 140.