INVESTIGATING THE PROCESS OF NATURAL LANGUAGE COMMUNICATION

A STATUS REPORT

Technical Note 165

July 1978

By: Ann E. Robinson
Artificial Intelligence Center
Computer Science and Technology Division

SRI Project 5844

The work reported herein was supported by the National Science Foundation under Grant No. MCS 76-22004.

ABSTRACT

This paper provides an overview of an ongoing research program in natural language communication, indicating its status as of June, 1978, and its short term goals. This research seeks to identify and computationally formalize the knowledge and processes needed for participation in natural language dialogs about ongoing tasks. The paper describes (1) the knowledge embodied in an existing system that interprets utterances in such dialogs, (2) the formalisms developed for encoding this knowledge, and (3) the framework in which the knowledge is combined and coordinated during the interpretation process. The paper also indicates anticipated extensions that will lead to refinements of interpretations. These extensions include the concept of modality, the use of the focus and goals of the dialog in the identification of the referents of pronouns, and the use of knowledge about the goals of the other dialog participant in the interpretation of utterances.

I SCOPE OF THE RESEARCH EFFORT

This paper provides an overview of the current status and short-term goals of an ongoing research program.* The current research builds on previous research in speech understanding performed at SRI from 1971 to 1976(Walker 1978). Detailed descriptions of particular aspects of our research are available in other publications(Grosz 1977a, 1977b, 1977c, 1978; Hendrix 1977, 1978; A. Robinson 1978; J. Robinson 1978a, 1978b).

The objective of our research is to increase our understanding of natural language communication by considering the underlying processes used to interpret and produce utterances (sentences or sentence fragments) exchanged during the course of a dialog. This research stands at the juncture of two disciplines: linguistics, which studies the units, structure, and meaning of natural languages, and computer science, which studies formal processes (algorithms) and their associated data structures and control mechanisms.

Linguists have generally concentrated their studies more on the form and meaning of individual words and sentences than on the processes through which information about some subject area is communicated. In contrast, our approach is to view natural language communication as an activity performed by two or more cooperating parties (either humans or intelligent computer systems). From this perspective, the focus of research is on the processes underlying participation in dialog rather than on the structure of utterances per se. The communication activity, including the subactivities of generating and interpreting (or "understanding") utterances, requires that dialog participants have knowledge not only about the structure of language; they must also be knowledgeable about the subject matter under discussion and about the processes through which concepts in the subject area may be expressed by means of words, phrases, and sentences of English.

----

* This research is currently funded under a three-year NSF Continuing Research Grant No. MCS76-22004. I would like to thank Gary Hendrix for his invaluable assistance in organizing and writing this paper.

To formalize knowledge about linguistic processes, our approach is to draw from (and ultimately to contribute to) computer science. Computer science seeks to provide formalisms for stating knowledge about algorithmic processes in precise terms. Because computer formalisms may be executed, they may be tested to determine whether they adequately encode the information needed to perform a given process. Our approach to the study of language is to use computer formalisms to express the knowledge needed to generate and interpret natural language utterances. But encoding information about natural language processing in computer systems introduces some interesting problems for computer science. Whereas convenient computer formalisms exist for representing the data and procedures involved in arithmetic operations, and research in artificial intelligence has contributed a number of formalisms for the manipulation of symbolic information, new formalisms must be developed to express such concepts as the focus of attention of language users, discourse structure, coherence, and intentionality--all of which play major roles in the analysis of natural language dialog.

Within the general framework of using computer formalisms as a means for investigating the communication processes, we are concentrating our attention on the study of dialogs between two parties cooperating to accomplish a shared task--the assembly or disassembly of an electromechanical device. We have chosen to study dialogs because they provide an opportunity to examine both the interpretation and production aspects of the communication activity. Moreover, task-oriented dialogs take place in a dynamic environment in which the progression of the task provides an ever-changing context. Most other current research in computational linguistics analyzes individual questions or statements within a static environment (as in data retrieval), or considers narratives rather than dialogs (as in story understanding).

We have collected a number of task-oriented dialogs between humans and have analyzed them to gain insights into the underlying communication processes. Some of the knowledge needed to participate in dialogs similar to those collected has been formalized in a computer system. This system acts as a dialog participant rather than as a nonparticipating observer. Should the ability of this system ever approach human proficiency, it would have numerous practical applications; however, it is not being developed with this as a primary goal. Rather, its development is an aid in determining what types of knowledge are needed for participation in task-oriented dialogs, for precise formalization of that knowledge, and for testing that formalization through execution of the computer system.

II PROGRESS TO DATE

Our activities center principally on the development of computational formalisms for representing increasingly significant subsets of the knowledge needed to participate in task-oriented dialogs. We are investigating the use of language from a procedural perspective and are developing an understanding of the dynamic interrelationships existing among the various types of information (including information about the subject domain and the structure of English) that contribute to the communication process. Simultaneously, these activities advance the development of new computer science techniques for representing and using knowledge about natural language and the subject matter under discussion. In the paragraphs below we describe a number of specific achievements and indicate their contributions to the total effort.

A. Formalizing Knowledge About the Subject Domain

A prerequisite for participation in the communication process is an ability to encode and utilize information about some subject area. We have therefore developed a representational formalism for encoding information about a dynamic environment and have used it to embody knowledge about objects, actors, relationships, and events relevant to the task of assembling an air compressor. The resulting domain model supports participation in dialogs pertaining to this subject matter by providing a semantic framework for the interpretation of utterances against the background of a changing environment, and by providing the subject-dependent information needed to respond to domain-specific questions.

The representation formalism is based on a network data structure that simultaneously encodes information and provides an index for accessing the information through semantic associations. A key feature of the structure is that the network may be partitioned into a number of overlapping subnetworks. These subnetworks, or spaces, may be used to delineate the scope of quantifiers, to establish contexts that specify the state-of-the-world situation at a given time, or to highlight those concepts that are in the current 'focus of attention' of the dialog participants.*

The network allows expression of both well-formed formulas based on first-order logic and, through newly developed formalisms called process models, information about processes in the assembly domain. The constructs from logic are used to describe the objects of the assembly domain, to indicate how they fall into hierarchical classes (a wrench is a tool, tools are physical objects), to indicate the actual and potential relationships among objects (the pump is on the table, one physical object may be on another physical object), as well as to indicate those relationships and events that existed in the past or are currently in progress.

Process models also indicate how various types of events may be hierarchically decomposed into subprocesses. They indicate the types of actors and objects that participate in an event, the preconditions for its enactment, its effects, and the alternative sequences of substeps that may be followed to accomplish it.

To derive information from the network data structures, we have developed procedures that embody general knowledge about logical

----

* The partitioned network formalisms developed at SRI are being used by researchers at other institutions, including the University of Toronto(Cohen 1976), the University of Texas (Simmons 1973), and Bolt Beranek and Newman (Woods 1975).

deduction and specific knowledge about how to manipulate process models. These procedures are used to retrieve or deduce answers to questions (which might be posed either in the process of interpreting an utterance or by the other participant in a dialog) and to assimilate into the model new information about the domain communicated by the user (e.g., the user may indicate that certain events have transpired or that new conditions have arisen).

An important aspect of participating in a dialog about a concurrent assembly operation is keeping an up-to-date account of assembly progress. The record of progress provides a context for interpreting utterances and responding to questions. At any given moment it is necessary to know what assembly events have already taken place (and in what sequence), what events are in progress, and what events can be initiated next. The establishment of a current point of reference provides a perspective on the state-of-the-world situation and a basis for determining which task steps appear most reasonable to perform next.

We have recently developed automatic procedures for following the course of simplified dialogs by extracting clues from the user's utterances regarding the progress of ongoing tasks. Because the linguistic structures employed by the user rely heavily on the situational context, the ability to keep an up-to-date account of the progress of the assembly and to assimilate changes into the model is of essential importance in dialog participation; consequently, its realization in an operational system may be considered one of our most significant achievements.

B. Formalizing Knowledge About Language

Our formalization of linguistic knowledge about the communication process may be divided into three major areas: (1) maintaining contexts for the interpretation and generation of utterances, (2) interpreting utterances, and (3) responding to utterances (including question-answering and the subsequent transformation of those replies into English). These areas are discussed in subsections below.

1. Contexts for Interpreting Utterances

Perhaps the most important achievement of our research to date has been the establishment of a framework within which to characterize language as it is used in a changing context. An utterance from a dialog cannot be meaningfully interpreted in isolation, but must be considered within the context in which it was produced. Interpretations are influenced both by the current state-of-the-world situation (and the history of past states) and by previous utterances in the dialog. State-of-the-world knowledge is essential, for example, in determining the referents of such phrases as "the tool in your hand". We call this particular aspect of the total context the task context. The importance of the prior discourse, which we call the discourse context, is illustrated by the pair of utterances

Speaker 1: Why did John take the pump apart?
Speaker 2: He did it to fix it.

in which the interpretation of the second utterance must take into account the context established by the first. In particular, "he" refers to John, "did it" refers to the disassembly task, and the second "it" refers to the pump.

We have created methods for representing information about both dialog context and task context. We have studied how this information is used to identify referents of definite noun phrases* in context and have formalized many of our concepts computationally.

The major aspect of context for which we have identified and developed a representation is called focus. At each instant, the focus of a dialog highlights a small collection of concepts which are directly or indirectly being discussed. Viewed as a process, focusing is an activity engaged in by dialog participants in which they concentrate

----

* Definite noun phrases (such as "the wrench" and "this pump") begin with a definite article or demonstrative adjective (such as "the" and "this"), as opposed to an indefinite article (such as "a") and are used to refer to particular objects. For example, consider the sentences:
"Shortly after John read War and Peace, Tom read the book too."
"Shortly after John read War and Peace, Tom read a book too."
"The book" refers to War and Peace, but "a book" does not indicate which particular book Tom read.

their attention on a subset of the information they share in common (ignoring the rest). By highlighting the subset of known concepts that are currently most relevant to them, language users create a localized frame of reference within which objects of significance to the dialog may be denoted simply and unambiguously.

We are currently working on three major subproblems associated with focus: (1) developing computational structures to represent focus, (2) determining from a dialog what entities are being focused upon, and (3) determining shifts in focus as a dialog progresses (i.e., perceiving the point of transition from one discussion topic to another). We have made substantial progress on the first two problems, but much remains to be done on all three.

C. The Process of Interpreting Utterances

To provide formalisms enabling us to specify the interactions of many types of domain and linguistic knowledge during the interpretation process, we have created a 'language definition system' called DIAMOND. DIAMOND may be viewed in broad outline as both a sophisticated programming language and an associated execution system. As a programming language, DIAMOND allows us to describe formally how utterances may be interpreted. The interpretation process is performed in multiple phases under control of the DIAMOND executive and in accordance with the specifications of the language definition formalisms. DIAMOND's flexibility allows us to experiment with different numbers of phases and with various sequences for applying various types of knowledge. Currently, we are working with three basic interpretation phases:

(1) The first phase is an analysis based on syntactic information about words and how they combine into phrases (i.e., words -> phrases -> larger phrases -> sentences). This phase produces a complex data structure (a generalization of a syntax tree) that reflects the decomposition of the input utterance into its component phrases and associates a number

of attributes with each phrase. This data structure serves as a guide for other phases of interpretation.

(2) Working from the data structure generated by the first phase, the second phase applies various knowledge resources to establish a relationship between phrases identified during syntactic analysis and <u>descriptions</u> of objects in the domain model. These knowledge resources must indicate the correspondences that exist between individual words and concepts in the domain model and the ways in which multiple simple concepts may be combined to form descriptions of more complex objects. (E.g., "box-end" names a shape, "wrench" names some member in the class WRENCHES, "box-end wrench" describes any member of WRENCHES that has a box-end shape).

(3) In the third phase, information about the current context is used to refine the interpretation. For example, phase 2 merely produces a description (in terms of the domain model) of a box-end wrench as its analysis of "the box-end wrench". In phase 3 "the box-end wrench" is associated with a particular object in the model. The individual object selected as the referent of the phrase "the box-end wrench" will differ according to the context. Moreover, knowledge about the domain and the strength of quantifiers is used for establishing the scope of quantifiers in phase 3.

Using DIAMOND, we have encoded and tested a general and linguistically-motivated set of syntax rules that cover a wide range of English constructions. We have successfully used these rules to parse (but not to assign interpretations to) sentences not only about the assembly task, but also about a retrieval task for a data base.*

----

* Supported by the Defense Advanced Research Projects Agency under Contract No. N00039-78-C-0060, with the Naval Electronics Systems Command.

Sample utterances covered by this syntax include the following:

What size wrench should I use to tighten the bolts?
The wrench that I used before is not on the table.
I've installed the brace, pulley and Woodruff key.
How many bolts of each size should I be using?

We have designed and implemented a set of 'semantic composition' procedures for use in the second phase of interpretation. Formalizing our knowledge of the relationship between words/syntactic structures in the language and concepts in the domain model, these procedures can associate proper names, such as "John", with individual constants in the model, such as JOHN. They can map common nouns, such as "wheel", into restricted free variables, such as "X where WHEEL(X)". They can map verb-like concepts, such as "removed" or "green-colored", into predicates in the model. In addition, these procedures can use knowledge about the domain to reject as meaningless such phrases as "the wheel removed John" (although the same string of words can be used in such meaningful constructions as "with the wheel removed John could fix it"). They can create network models to represent meaningful phrases, such as "John removed the wheel" by combining information about the individual words, the syntax that relates them, and the relationships potentially extant in the assembly domain. Semantic composition routines are much more difficult to formalize than syntax rules. Consequently, we do not yet have composition procedures covering all the constructions that can be analyzed syntactically.

1. <u>Responding to Utterances</u>

Part of communication is determining the appropriate action to take once an utterance has been interpreted. This means much more than simply answering questions that are posed by the other party (although we do have provisions for answering questions). Each input must be analyzed for clues regarding changes in the user's focus of attention, as well as changes in the progress of ongoing tasks. A major part of our research effort is to discover the mechanisms through which dialog participants convey such clues.

To respond to queries we have evolved procedures which answer certain types of questions by making deductions from the information in the domain model. A network-to-English translation system developed in a previous project(Slocum 1978) produces English replies to questions, using simple syntactic constructions.

D. State of Implementation and Testing

By implementing new computer formalisms and encoding knowledge about the communication process in them, we have created an operational computer system whose performance tests the adequacy of our theories. Section V contains a segment of a dialog in which the system has participated. We characterize below the scope of the system's current capabilities.

A unique aspect of our implementation is that it operates in and is responsive to a dynamically changing environment. As utterances indicate that progress has been made in the task and/or that the user is shifting the focus of the dialog, the change in context is reflected in the system's model of the situation. This influences the way in which subsequent utterances are interpreted. Changes can be indicated either by statements which directly specify a change ("I have attached the pump") or less directly by utterances about particular objects or substeps ("Where is the Allen wrench?").

In addition to interpreting statements which denote changes in the task state, the system can interpret and respond to questions about the current state-of-the-world situation ("Where is the box-end wrench?"), past states of the world ("Was the pump on the table before it was attached?"), how to perform a step in the task ("How do I install the pulley?"), and what step should be performed next ("What should I do now?"). Similarly, the system can deal with questions about participants in and results of previous steps ("Who attached the pump?", "Have the bolts been tightened?").

A history of the previous events and states of the world is maintained from which to derive responses to queries about the past and to help interpret phrases such as "the wrench that was used to tighten the motor bolts". To process questions about the assembly task, the steps necessary for accomplishinging it are encoded and can be used to deduce the effect of performing an action. The same encoding is used to follow the dynamics of the dialog. That is, as a task is performed, the system determines the new state-of-the-world situation, which includes the results of intermediate steps. The computer representation of the dialog's focus is then updated to reflect this revised situation, so that references to specific entities can be correctly identified in subsequent inputs.

III RESEARCH PLANS

During the following year we plan to refine, extend, and further test those knowledge resources for which we have already developed formalisms. By introducing new types of information and refining those we have already encoded, we will both broaden the range of utterances which can be interpreted and derive more information from the utterances currently processed. Specific areas to which we plan to devote our efforts include: (1) significantly refining the model of the assembly domain, (2) developing procedures for linking complex syntactic constructions to both objects and relations in the domain model, (3) using focus to determine referents of pronouns, (4) studying the ways in which the interpretation of references to events (as opposed to objects) is influenced by the context, (5) studying modal verbs, and (6) considering how knowledge about the user's goals and beliefs may facilitate interpretation of utterances beyond their strictly literal meanings.

A. Refining the Model

In refining the model, we plan to extend the task model to cover a larger portion of the assembly task, refine our descriptions of objects and their interrelationships in the domain, and to add other processes which are not directly associated with the task (e.g., "plug in", "find", "move an object"). This refined model will permit more sophisticated testing of our representational formalisms and enable us to test our language processing algorithms with language structures covering a wider range of concepts directly and indirectly related to the assembly task.

B. Linking Complex Syntactic Structures to Entities in the Domain Model

At present we have a syntactic description of a broad range of English constructions. However, computational descriptions for advanced phases of interpretation exist for only a subset of the utterances analyzed syntactically. For example, sentences with embedded clauses, such as "I am using the belt which was in the toolbox" are currently analyzed into phrase units, but are not otherwise being interpreted. In particular, there is no link established between these types of sentences and the objects and relationships in the domain model. Over the next year we plan to extend the computational descriptions of the interpretation process to cover such embedded-clause utterances. This will expand the range of English for which we have a formal characterization and increase our knowledge about the ways in which the interpretation of constituents of complex sentences is related to the interpretation of the sentence as a whole.

C. Using Focus to Determine Referents of Pronouns

Our current algorithm for determining referents of pronouns does not use information about which items are in focus, but instead searches back through previous sentences in the dialog for the nearest noun phrase that agrees with the pronoun in gender and number. This method works in many cases but fails in others. We plan to experiment with the theory that pronouns are used to refer to those concepts that are most prominently in focus rather than those most recently mentioned. To be prominently in focus, a concept need not have been mentioned recently (e.g., in some of the dialogs that we collected, the pronoun "it" was used to denote a concept--the air compressor--that had not been referred to directly for over 30 minutes) or explicitly (e.g., "The Joneses are coming to dinner and are bringing HER mother.").

D. Interpreting References to Events in Context

Interpreting references to events in context is the verb phrase analogue to the contextual identification of noun phrase referents. We believe that many of the methods developed by us for noun phrase resolution will be applicable to this problem. As the following examples show, there are many ways of referring to an event, ranging from highly specific (as in the first example) to short utterances composed of pronouns and pro-verbs (as in the last example).

> I positioned the pump on the platform.
> I put the pump on the platform.
> I put it there.
> I did it.

As references to events become less explicit, interpretation is increasingly dependent upon the surrounding contexts of task and dialog. We plan to investigate the different ways in which events are referred to and the role of the context in interpreting references to them. We will characterize different types of references to events, describe the complexities of reference-context interaction, and formalize our analysis by developing algorithms for interpreting references to events within the framework of our operational system.

E. Modal Verbs

We plan to extend our computational framework to include the concepts of modality, which are expressed in English by such modal verbs as "can", "could", "should", "will", and "must". Modal verbs are of interest because their inclusion in our general framework would add a new dimension to the types of relationships under investigation. Most ordinary verbs (e.g., "assemble", "move", "tighten") are used to refer to events involving actors and objects. But modals operate at a meta-level, referring to events per se and implying the possibility or necessity of such events. Modals also project predictively into the future and can indicate whether a given action is consistent with the objectives of performing a certain set of tasks. We believe that the dynamic nature of our domain model provides a powerful new framework within which to study these verbs. For example, the task models already indicate what should be done next to progress further in accomplishment of the task. We plan to add new information to indicate which events are capable of occurring and to extend our deduction mechanisms to reason about possible future conditions that might or will unfold.

F. Interpretation of Utterances Beyond Their Strictly Literal Meaning

Often the literal interpretation of an utterance captures only a portion of the total meaning intended. For example, an utterance such as "Is there a wrench?" can be intended to mean "If there is a wrench, tell me where it is and how to find it" instead of merely "Tell me if there is a wrench". Given this interpretation, a useful and appropriate answer would be "Yes, the wrench is in the box in the tool cabinet" rather than "Yes, there is a wrench." To interpret the preceding example properly, the hearer must know that the speaker has as a goal the performance of some action that requires a wrench, that the speaker does not know whether a wrench is nearby, that the speaker could accomplish the goal if the location of the wrench were known, that the hearer is expected to help the speaker, if possible, and that the hearer can help by informing the speaker as to the location of the wrench.

Such reasoning requires augmenting the knowledge we have now represented about language with knowledge which includes information about the speaker's goals, both immediate and long-term, and his underlying beliefs. Over the next year, we will begin investigating the ways in which such knowledge can be used in the interpretation process and, in particular, how it can be incorporated into our computational framework.

IV PROJECTED STATUS

Our current research will result in development of formal descriptions of how certain types of knowledge about language and about the world are used by dialog participants. These descriptions will be presented as computational formalisms, many of which will have been implemented and tested in a computer system.

These formalisms include:

* Formalisms for encoding various types of knowledge about the domain of discourse (encompassing objects, states, and events and their participants).
* Formalisms for encoding various types of knowledge about linguistic processes (including dialog focus, syntactic analysis, and semantic association).
* An overall framework for describing how multiple types of knowledge interact in the communication process.

Encoded in the formalisms will be the following types of information, each of which plays an important role in the communication process:

* The vocabulary of a sizable portion of English relevant to the subject domain.
* The objects, relationships, states, and processes in the subject domain.
* How words/phrases may be combined to form larger linguistic units.
* The relationship between words and phrases in English and concepts in the domain model (including phrases involving tense and modality).

* The focus of a dialog at any point in time.
* How the dialog focus facilitates the identification of entities referred to by definite noun phrases, verb phrases, and pronouns.
* How the focus of a dialog shifts as the task and dialog progress.
* How to generate responses to utterances, including deducing answers to questions about objects, events, and processes (past, present, future, and hypothetical).
* How to monitor changes in the state-of-the-world situation.
* Preliminary descriptions of how the goals and beliefs of dialog participants influence the interpretation of utterances.

Viewed as a whole, this work will provide a coherent description of a major portion of the communication process underlying task-oriented dialogs. This description will provide a valuable basis for future research. Some important and potentially very interesting areas of future research are:

* Formalization of the way in which the goals and beliefs of dialog participants influence the interpretation of utterances.
* Formalization of the additional elements contributing to context--particularly shared visual context and implicit social conventions.
* Production (as opposed to interpretation) of utterances that take into account what one's dialog partners are thought to know or believe.

V Sample Dialog with our Computer System

We have created an operational computer system whose participation in an actual dialog with a human is illustrated below. The input to the system is preceded by a '>'; the rest of the dialog was produced by the computer system and illustrates how it monitors the performance of the task. The numbers following the input are used in the commentary that follows.

```
>HOW DO I INSTALL THE PUMP?                    (1)

Attach the pump to the platform.
   then in either order:
     Install the aftercooler elbow on the pump.
        then
     Install the aftercooler on the pump.
   and
     Install the brace on the pump.
        then
     Install the pulley on the pump.


>WHERE IS THE BOX-END WRENCH?                  (2)
The box-end wrench W4 is on the table.

>I AM TIGHTENING THE BOLTS WITH THE WRENCH.    (3)
OK
                The following has been assumed:
                    You put the pump at the platform.

>I'VE ATTACHED THE PUMP.                       (4)
OK


>WHAT SHOULD I DO NOW?                         (5)

Install the aftercooler elbow on the pump.
     or
Install the brace on the pump.


>THE BRACE HAS BEEN INSTALLED.                 (6)
OK
```

>WHERE IS THE WOODRUFF KEY? (7)
The Woodruff key is on the table.

>I HAVE INSTALLED IT. (8)
OK

>THE PULLEY HAS BEEN INSTALLED. (9)
OK

The following has been assumed:
You attached the pulley to the pump.
You put the pulley at the pump.
You fastened the pulley to the pump with the screws.

This dialog segment covers part of the pump installment subtask of the compressor assembly. The reply to the first question indicates the steps to be performed. Not only is the progress of the task followed, as is indicated by the computer generated commentary on the right of the dialog, but the focus of the dialog shifts accordingly. Initially, the focus includes the user, the compressor, the pump, and the table. As an object or step is mentioned, the focus shifts to make that object or the entities associated with that step the primary focus. Thus, the second utterance puts the box-end wrench into focus. There are several wrenches in the model but only one box-end wrench, so it can be identified from the phrase "the box-end wrench" but not from the phrase "the wrench" in isolation. Once the box-end wrench is in focus, after the second utterance, subsequent references to "the wrench" will be sufficient to identify it as long as it remains in focus. The fourth and sixth utterances put the pump and platform, and then the brace and pump into primary focus because they are associated directly with the steps mentioned in the task. The Woodruff key in the seventh utterance is employed in installing the pulley, but is associated with a substep. However, because it is indirectly associated with the installation task--the next step to be done--it is in 'implicit focus' and can still be identified by the system and then moved into primary focus. The last utterance shows completion of the 'install pulley' task.

VI REFERENCES

Cohen, Philip R., and Perrault, C. R. "Preliminaries for a Computer Model of Conversation." In Proceedings of the First CSCSI/SCEIO National Conference, Vancouver, British Columbia, Canada, August 25-27, 1976, pp. 102-111.

Grosz, Barbara J. The Representation and Use of Focus in Dialogue Understanding. Technical Note No. 151, SRI International, Menlo Park, California, 1977.

Grosz, Barbara J. The Representation and Use of Focus in a System for Understanding Dialogs. Proceedings of the Fifth International Joint Conference on Artificial Intelligence, Pp. 67-76, Cambridge, Massachusetts, 22-25 August 1977.

Grosz, Barbara J., Hendrix, Gary G., and Robinson, Ann E. Using Process Knowledge in Understanding Task-Oriented Dialogs. Proceedings of the Fifth International Joint Conference on Artificial Intelligence, P. 90. Cambridge, Massachusetts, 22-25 August 1977.

Grosz, Barbara J. Focusing in Dialog. Proceedings of TINLAP-2, Urbana, Illinois, 24-26 July, 1978.

Hendrix, Gary G. Some General Comments on Semantic Networks. Panel on Knowledge Representation, Proceedings of the Fifth International Joint Conference on Artificial Intelligence, Pp. 984-985, Cambridge, Massachusetts, 22-25 August 1977.

Hendrix, Gary G. Encoding Knowledge in Partitioned Networks. In Associative Networks - The Representation and Use of Knowledge in Computers. Edited by Nicholas V. Findler, Academic Press, New York, New York, forthcoming.

Robinson, Ann E. Speech, Language and Artificial Intelligence. Proceedings of the Seventh International Conference on Computational Linguistics, Bergen, Norway, 14-18 August 1978.

Robinson, Jane J., and Hobbs, Jerry R. Why do You Ask? Proceedings of Linguistic Society of America Summer Meeting, University of Illinois, Urbana-Champaign, Illinois, 28-30 July, 1978.

Robinson, Jane J. Pointed Questions and Useful Answers. To appear in Proceedings of the Seventh International Conference on Computational Linguistics, Bergen, Norway, 14-18 August 1978.

Simmons, Robert F. "Semantic Networks: Their Computation and Use for Understanding English Sentences." in Computer Models of Thought and Language. Roger C. Schank and Kenneth M. Colby (eds.), Freeman, San Francisco, 1973, pp. 63-113.

Slocum, Jonathan, "Generating a Verbal Response." in Understanding Spoken Language. Elsevier North-Holland, New York, 1978.

Walker, Donald E., (ed) Understanding Spoken Language. Elsevier North-Holland, New York, 1978.

Woods, William A. "What's in a Link: Foundations for Semantic Networks." In Representation and Understanding. Daniel G. Bobrow and Allen M. Collins, (eds.), Academic Press, New York, 1975, pp 35-82.