A COMPREHENSION MODEL FOR HUMAN DIALOCVE

William C. Mann, James A. Moore, James A. Levin

USC/Information Sciences Institute
Marina Del Rey, California 90291

ABSTRACT

The comprehension of dialogue is an important concern for those interested in natural language processing for several reasons: dialogue gives particularly good access to human communication phenomena, it is less contrived than authored text, and human dialogue provides useful analogies for improving man-machine communication. In naturally occurring dialogues, the goals of the participants play a key role in structuring their language interactions. People know how dialogue is used to achieve goals, and they use this knowledge to comprehend what they hear.

We have represented this knowledge in structures called Dialogue-games, which differ from other multisentential knowledge structures by representing knowledge about how language is used to pursue goals rather than the structure of the content being conveyed. These Dialogue-games, which govern such activities as helping, information seeking, giving instructions, announcing, and testing someone's knowledge, are part of a strongly goal-oriented model of language comprehension. This model is composed of autonomous processes which operate on shared memories. Each of the four memories (a long-term memory and a short-term memory for each of two participants) is a collection of predicate expressions representing the individual's fixed knowledge and current awareness.

An important part of this research has focused on the problem of model evaluation. Benchmarks for use in evaluating model performance have been developed, creating opportunities for empirical validation of the effectiveness of natural language process models.

The modeling effort described here has led to progress on several problems of natural language comprehension - it provides an explanation for topic changes, serves as a basis for limiting inferences, accounts for the comprehension of a number of different kinds of implicit communication, and suggests ways that formal man-machine communication could be improved. It thus illustrates the utility of an overall goal-oriented view of language.

The research reported in this paper was supported in part by the Personnel and Training Research Programs of the Office of Naval Research, Contract N00014-75-C-0710, under direction of the Cybernetics Technology Office of the Advanced Research Projects Agency.

STATEMENT OF THE PROBLEM

The principal goal of this research has been to *create a model of human communication* at an appropriate level of detail to benefit man-machine communication design (Mann, 1977). It has included three closely related investigations:

1. The study of naturally occurring language to discover regularities of usage and to determine how these regularities contribute to the communicative aspects of language.

2. The represention of the understanding of these regularities as data structures and process models.

3. The establishment of standards by which the model's performance can be compared with that of humans on closely related tasks.

It has been further limited to:

1. only modeling the *receptive* aspects of communication,

2. only examining *dialogue,*

3. only modeling dialogue conducted over a *restricted medium* so that there is no visual or intonational communication, (which would not be captured in the transcript.)

A model, in this research, is a set of interacting processes which takes a transcript of an actual dialogue, turn by turn, as input, and repeatedly updates a memory (called a Workspace) whose content represents the current awareness of the receiver (hearer) of the most recent turn (Figure 1). The model is evaluated by comparing the changes of the two Workspaces with human judgments about the input transcript; judgments for which there are corresponding Workspace changes are model successes.

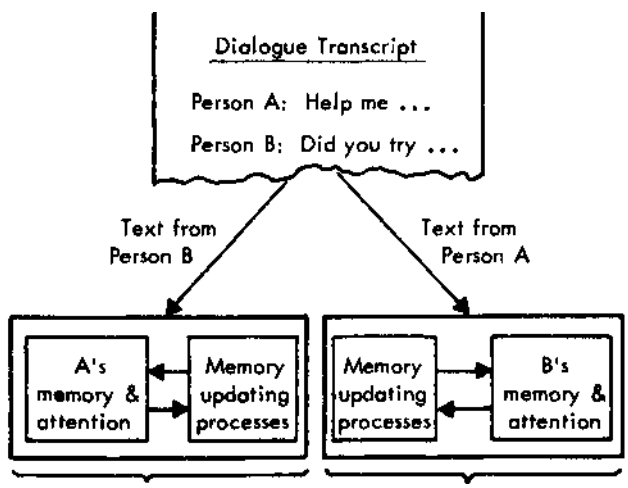

Figure 1. Simplified information flow of the Dialogue Modeling System.

The research has dealt not only with the traditional problem of model creation, but also with the problem of model success identification, especially of eliciting suitable human judgments so that model successes become significant.

PAST RESEARCH ON LANGUAGE COMPREHENSION

There is a great diversity of research potentially relevant to modeling language comprehension, at levels of detail ranging from units smaller than words to communication behavior taken in aggregales of many years, many individuals and many thousands of words. The level of detail of interest here contains individual transactions, accomplished in a short time (of the order of seconds or minutes), and unit sizes ranging from morphemes to short conversations. Within these limits there is still an abundance of work and of viewpoints. We take the most relevant to be those of theoretical linguistics (including syntax, semantics and higher), computer science (mainly artificial intelligence), and cognitive psychology. These three provide a number of constructive suggestions on model form as well as constraints on performance.

We are particularly interested in issues of how symbols affect the intended receiver, how they change his state. How does the state-change produced by a sentence depend on previous sentences, on understandings about how language is used, or on shared knowledge of the world? There is a widespread recognition that there are regularities of language use that ordinarily span multiple sentences. Narrative stories have been studied by several groups to identify their recurrent knowledge structures. A specific model for the form of this multisentential knowledge is the "story schema", organized within a story grammar (Rumelhart, 1975). This model has been supported by the results of story recalls (Rumelhart, 1975; Thorndyke, 1977). Other similar kinds of theoretical constructs for organizing multiple sentences of stories have been proposed called: "frames" (Minsky, 1975; Charniak, 1975), "scripts" (Schank *&* Abelson, 1975), and "commonsense algorithms" (Rieger, 1975).

To account for the conduct and comprehension of dialogues, multisentential knowledge units have also been proposed by linguists and sociolinguists to explain certain kinds of regularities observed in naturally occurring dialogues. These regularities have been called "rules" by Labov *&* Fanshel (1974) and "sequences" by Sacks, Schegloff, *&* Jefferson (1974).

Once these multisentential knowledge units are evoked, they serve as a basis for comprehending the successive inputs. This is achieved by generating expectations and by providing a framework for integrating the comprehension of an utterance with that of its predecessors. Recently, we proposed (Levin *&* Moore, 1976, 1977) multisentential knowledge units that are specified primarily by the speaker's and hearer's goals. These differ from the other proposed multisentential units by representing knowledge about how language is used to pursue goals rather than representing the structure of the content being conveyed. These goal-oriented units, which we call Dialogue-games*, specify the kinds of language interactions in which people engage, rather than the specific content of these interactions.

Other closely related research by Barbara Grosz has investigated the structure of task oriented dialogues between expert and novice. (Deutsch, 1974; Grosz 1976j Grosz 1977) She found that significant parts of the dialogue structure were derived from the goal structure of the shared task. This structure was active *in* resolving references, causing some distant potential reference objects to be preferred to recent ones. It also affected modes of acknowledgement, interruption structure, noun phrase reference and coherence under segment deletion.

An important problem facing researchers in language comprehension is posed by sentences with which the speaker performs what philosophers of language have called "indirect speech acts" (Searle, 1975). The direct comprehension of these sentences fails to derive the main communicative effect. For example, declarative sentences can be used to seek information ("I need to know your social security number."); questions can be used to convey information ("Did you know that John and Harriet got married[9]") or to request an action ("Could you pass the salt?"). These kinds of utterances, which have been extensively analyzed by philosophers of language (Austin, 1962; Searle, 1969, 1975; Grice, 1975), are not handled satisfactorily by any of the current theories of the direct comprehension of language. However, these indirect language usages are widespread in naturally occurring language—even two year old children can comprehend indirect requests for action almost as well as direct requests (Shatz, 1975).

One theory proposed to account for these indirect uses of language is based on the concept of "conversational postulates" (Grice, 1975; Gordon *&* Lakoff, 1971). If the attempt at direct comprehension of an utterance produces an implausible result, then the indirect meaning is derived using these postulates.

In general, this approach to indirect speech acts is inference-based, depending on the application of conversational rules to infer the indirect meaning from the direct meaning and the context. A different approach has been proposed by Labov *Si* Fanshel (1974) and by Levin *&* Moore (1976). Multisentential knowledge, organizing a segment of language interaction, can form the basis for deriving the indirect effect of utterance within the segment. For example, a multisentential structure for an information-seeking interaction can supply the appropriate context for interpreting the subsequent utterances to seek and then supply information. The inference-based approach requires one set of conversational rules for information requests, a different set of rules for answers to these requests, and a way to tie these two rule sets together. The Dialogue-game model postulates that there is but one knowledge structure for this kind of interaction, and leads to a model of three sets of cooperating processes: (I) processes for recognizing when this kind of interaction is proposed, (2) processes for using this knowledge to comprehend utterances within its scope, and (3) processes for identifying when the interaction is to be terminated.

* The term "Dialogue-game" was adopted by analogy from Wittgenstein's term "language game" (Wittgenstein, 1958). However, Dialogue-games represent knowledge people have about language as used to pursue goals, rather than Wittgenstein's more comprehensive notion. Although there are also similarities with other "games," the properties of Dialogue-games are only those described here. For example, they are not necessarily competitive or consciously pursued.

THE SHAPE OF THE THEORY

Our theory of human use of language has been strongly influenced by work in human problem solving (Newell & Simon, 1972), in which the behavior of a human is modeled as an information-processing system, having goals to pursue and selecting actions which tend to achieve the goals. We view humans as engaging in linguistic behavior in order to advance the state of certain of their goals. They decide to use language, they select (or accept) the other participant for a dialogue, they choose the details of linguistic expression — all with the expectation that some of their desired state specifications can thereby be realized. Furthermore, they break off an interaction either when the relevant goals have been satisfied, or when it becomes clear that they cannot be.

In this theory of language, a participant in a linguistic exchange views the other as an independent information processing system, with separate knowledge, goals, abilities and access to the world. A speaker has a range of potential changes he can effect in his listener, a corresponding collection of linguistic actions which may result in each such change, and some notion of the consequences of performing each of these. The speaker may view the hearer as a resource for information, a potential actor, or as an object to be molded into some desired state.

A dialogue involves two speakers, who alternate as hearers. By choosing to initiate or continue the exchange, a participant attempts to satisfy his own goals; in interpreting an utterance of his partner, each participant attempts to find the way in which that utterance serves the goals of his partner. Thus a dialogue continues because the participants can continue to see it as furthering their own goals. Likewise, when the dialogue no longer serves the goals of one of the participants, it is redirected to new goals or terminated.

This mechanism of joint interaction, via exchange of utterances, in pursuit of desired states, is useful for achieving certain related pairs of participants' goals (e.g., learning/teaching, buying/selling, getting help/giving help, ...). For many of these goal-pairs there are highly structured collections of knowledge, shared by the members of the language community. These collections specify such things as: 1) what characteristics an individual must have to engage in a dialogue of this sort, 2) how this dialogue is initiated, pursued and terminated, 3) what range of information can be communicated implicitly, and 4) under what circumstances the dialogue will "succeed" (serve the function for which it was initiated) and how this will be exhibited in the participants' behavior.

This characterization yields technical explanations of a number of interesting language phenomena.

- It models topic structure well. People can identify the beginnings and endings of topics in dialogue very reliably. Topics are taken up and dropped as part of taking up and dropping particular goals. Most goal changes are performed by changes in the set of games in effect in the course of conversation. The fact that topic endings can be detected and explained systematically is particularly interesting, since people often do not signal topic endings explicitly.

- It provides a principled basis for limiting the amount of inference included in the comprehension process, since completion of comprehension occurs when all of the text of a turn has been identified with pursuit of particular goals of the speaker.

- It explains a great deal of implicit communication, including both "indirect speech acts" and communication of preconditions of actions. (For example, the question "Can you help me send a message?" is not simply a question about one's ability to help. It communicates a request for help as an indirect speech act, and it communicates that the speaker wants to send a message by precondition communication.) Since implicit communication is accounted for, part of the surprising brevity of human communication is also accounted for in this way.

- It suggests the technical causes of persistent difficulties in man-machine communication. This formulation of communication and comprehension is in sharp contrast with typical practise in interactive man-machine communication. System interfaces *neVer* analyze why a particular command is given, why a particular parameter is called for or how the results of executing a command will be useful to the commander. The kinds of knowledge structures and processing that people use to facilitate and control their communication with each other are absent from man-machine interfaces. We see this as one of the dominant sources of the difficulties that people, especially computer-naive people, have in communicating with machines.

A DIALOGUE *EXAMPLE*

Figure 2 shows a dialogue which is representative of the materials we have analyzed in detail. The participants are a computer operator 0 and a user L of the TENEX timesharing system. The user initiated the dialogue by a LINK command, which causes his terminal and (O's terminal to both display any characters which would normally appear on either one. They then engage in dialogue by typing to each other. They are in separate locations and cannot see each other, and they do not know each other by name*

We would like the model to analyze this dialogue into parts such as those indicated by brackets on the right side in Figure 2. The entire interaction is an action-seeking interaction with L's principal goal being to reduce his use of disk space while retaining access to files currently stored on disk. A required step in his chosen method for achieving this goal is to assign a magnetic tape to L's group. The system operator, 0, has the right to do this assignment on request. This assignment action is being sought in pursuit of L's principal goal.

* Brackets [] indicate that a name has been altered for privacy. A dectape is a small magnetic tape suitable for off-line storage of files.

Action- seeking going

Turn 1: L
Is there a free dectape that can be assigned that we may release
some disk space?

Turn 2: 0
Yes.

Turn 3: L
Can you assign it to us then?

Action- seeking game

Turn 4: 0
[user-group name]? Information- see *game*

Turn 5: L
Yes.

Turn 6: 0
OKThat will be Dectape 0173.

Turn 7: L
O.K. Can you mount that now?

Act/o -seek/if game* - T J

Turn 8: 0
i'm sorry no dectape drives are unavailable right now

Turn 9: L
OkayJ Later. Bye.

Turn 10: 0
Bye.

Figure 2. Structure of example dialogue

Turn 1 initiates the action-seeking and also initiates a more specific information seeking interaction. The last phrase in turn 1 is relevant to the general action-seeking interaction since it reveals L's principal goal; it is irrelevant to the information-seeking interaction of turns 1 and 2.

Turn 3 appears on the surface to initiate information-seeking but actually functions as a request and so initiates action-seeking.

Turn 4 could be paraphrased "Are you a member of [user-group name]?" It indicates that 0 is beginning to perform the assignment. Turn 5 is an answer to turn 4.

Turn 6 indicates completion of the action, and the "O.K."* in turn 7 indicates acceptance of the termination of that action-seeking interaction. In turn 7 a new action seeking interaction is initiated in pursuit of I's goal of releasing space. The action is refused with a reason in turn 8, and acceptance of termination of the action-seeking interaction is given in turn 9. The overall interaction then terminates, since both parties recognize that the main goal cannot be further pursued at the moment.

We would like a formal account for these sorts of perceptions, specified well enough so that processes can be defined to *carry* out the operations specified in the account.

THE DILOCUECAME MODEL

This section describes the Dialogue-game Model at its current state of development. It starts with a brief overview of dialogue and how it is regulated, then describes the dominant knowledge structures which guide the model, and finally describes a set of processes which apply these knowledge structures to dialogue utterances, comprehending them.

Each participant in dialogue is simply pursuing his own goals of the moment. The two participants interact smoothly because the conventions of communication coordinate their goals and give them continuing reasons to speak and listen. These goals have a number of attributes which are not necessarily consequences of either human activity in general, or communication in particular, but which are nonetheless characteristic of human communication in the form of dialogue:

1. *They are cooperatively established.* There are bidding and acceptance activities that serve to introduce goals.

2. *They are mutually known.* Each party assumes Or comes to know goals of the other, and each interprets the entire dialogue relative to currently known goals.

3. *They are configured by convention.* Sets of goals for use in dialogue (and other language use as well) are tacitly known and employed by all competent speakers of the language.

4. *They are bilateral.* Each dialogue participant assumes goals complementary to those of his partner.

5. *They are uhiquitou*.* A hearer views the speaker as always having goals he is pursuing by speaking. Furthermore, the hearer recognizes and uses these goals as part of his understanding of the utterance.

An uninterrupted dialogue goes through three phases:

establishing goals,
pursuing goals,
decommitting from goals.

Typically this sequence is compounded and repeated several times in the course of a few minutes.

We have created knowledge structures to represent these conventions, and processes to apply the conventions to actual dialogues to comprehend them. Since the knowledge structures dominate all of the activity, they are described first.

Text is interpreted in this model by frequent modification of a -Workspace" which represents the attention or awareness of the listening party. The modifications are roughly cyclic:

1. A new item of text T is brought into attention through the Parser.

2. Interpretive consequences of T are developed in the Workspace by a variety of processes.

3. An expression E appears in the Workspace which relates T to an imputed goal G of the speaker of T.

This final expression E is of course a formal expression in the knowledge representation of the model. E represents the proposition (held by the hearer) that in uttering T, the speaker was performing an act in pursuit of G, a speaker's goal known to the hearer. Comprehension is equated with relating text to pursuit of speaker's goals.

To make an explicit account of dialogue in this way, we now describe the knowledge structures which represent those conventions that supply the goals for the participants to pursue. In particular, we will answer the following three questions:

1. What is the knowledge we are representing within the definition of a particular Dialogue-game?

2. How is this knowledge used to model the receptive acts of dialogue participants?

3. What sorts of processes does it take to perform the receptive acts specified by this model?

The Dialo gue- game Knowledge Structure

A Dialogue-game consists of three parts: a set of *Parameter*,* the collection of *Specification** that apply to these Parameters throughout the conduct of the game, and a partially ordered set of *Component** characterizing the dynamic aspects of the game. For the balance of this section, we will elaborate on these three parts and exemplify these with an example of the Helping-game.

Bidding and Acceptance are entry operations which people use to enter Dialogue-games. Bidding:

1. identifies the game,

2. indicates the bidder's interest in pursuing the game,

3. identifies the Parameter configuration intended.

Bidding is performed many different ways, often very briefly. It is typically the source of a great deal of implicit communication, since a brief bid can communicate all of the Parameters and their Specifications for the Dialogue-game being bid.

Acceptance is one of the typical responses to a Bid, and leads to pursuit of the game. Acceptance exhibits:

1. acknowledgement that a bid has been made,

2. recognition of the particular Dialogue-game and Parameters bid,

3. agreement to pursue the game,

4. assumption of the Acceptor's role in the Dialogue-game.

Acceptance is often implicit, especially in relatively informal dialogue. It can be indicated by statements of agreement or approval, or by beginning to pursue the game (i.e. attempts to satisfy the goals). Alternatives to acceptance include rejection, negotiation and ignoring.

Bidding and acceptance appear to be part of game entry for all of the Dialogue-games of ordinary adult dialogue. Bidding and acceptance can also lead to game termination. Games can also terminate by satisfaction of the principal goal, by unconditional failure of that goal or by interruption.

*Parameter**

Dialogue-games capture a certain collection of information, common across many dialogues. However, the individual participants involved, and the topic (but not the function) of the dialogue may vary freely over dialogues described by the same Dialogue-game. To represent this, each Dialoguo-game has a set of Parameters which assume specific values for each particular dialogue.

The dialogue types we have represented so far as Dialogue-games have required only these three Parameters: the two participants involved (called "Roles"), and the subject of the dialogue (called "Topic"). This Topic is currently unconstrained — any concept representable in memory.

*Parameter Specification**

One of the major aspects distinguishing various types of dialogues is the set of goals held by the participants. Another such aspect is pattern of known and unknown information of the participants. We have found that for each type of dialogue, there is a corresponding set of descriptions which must hold for the goal and knowledge states of the participants, vis-a-vis each other and the subject. Within the formalism of the Dialogue-game, these are called the Parameter Specifications, and are represented by a collection of predicates on the Parameters.

The requirement that these specifications be satisfied throughout the conduct of a game is used by the participants to: signal what game(s) they wish to conduct, recognize what game is being bid, decide how to respond to a bid, conduct the game once the bid is accepted and terminate the game when appropriate. These Specifications also provide the means with which to explain the implicit, but clearly successful, communication which accompanies almost all natural dialogue. Examples and discussions of these Specifications will accompany the example of the Helping-game, below.

*Component**

The Parameter Specifications represent aspects of a dialogue that remain constant throughout the course of a dialogue of that type. We have also found that certain aspects change in systematic ways; these are represented in Dialogue-games as Components. In the Dialogue-games we have developed so far, the Components have been represented as a set of participants' subgoals, partially ordered in time.

Once a game has been bid and accepted, the two participants each pursue the subgoals specified for their role by the Components of this game. These subgoals are often mutually complementary, each set facilitating the other. Furthermore, by the time the termination stage has been reached (subject to a few constraints), pursuit of the Component subgoals will have assured satisfaction of the higher, initial goals of the participants, in service of which the game was initiated in the first place.

The Il el ping-(fame, an Example

In this section, we exhibit a specific Dialogue-game: the H*el ping-game.* This game is presented in an informal representation, in order to emphasize the informational content, rather than the representational power of the formalism. In what follows, the italics indicates the information contained in the representation of this particular Dialogue-game; the intervening text is explanatory commentary. (A number of other Dialogue-games are described in similar detail in Levin & Moore (1977).)

The (annotated) *Helping-game*

Parameter: HELPEE, HELPER, and TASK.*

The HELPEE wants help from the HELPER. The TASK is some sort of a problem, otherwise unspecified.

Parameter Specifications:

HELPEE: want to perform TASK.*

HELPEE: want to he able to perform TASK.*

HELPEE: not able to perform TASK.

HELPEE: permitted to perform TASK.

HELPEE: n person.

These Specifications not only constrain who would qualify as filling the role of HELPEE, but also provide reliable information about the HELPEE, given that this individual is believed to be engaged in the Helping-game. So, if someone asks for help on a task, we can assume that he wants the task performed. He is regarded as "insincere" if he does not want the task performed and yet asks for help in performing it, and he implicitly communicates that he wants the task performed by asking for the help.

HELPER: want to help HELPEE perform TASK.*

HELPER: able to provide hi help.*

*HELPER: a per*on.*

So, in order to be a HELPER, an individual must be willing and able to provide the needed assistance. Since this Dialogue-game represents **hared* knowledge, the HELPEE knows these Specifications, and therefore will not bid the Helping-game to someone who is not likely to meet them. And similarly, no one who fails to meet these Specifications (and knows he fails) will accept a bid for the Helping-game with himself as HELPER.

Component of the Helping-game:*

There are three components; the first two constitute the "Diagnosis" phase to communicate what the problem is.

*/. HELPEE wants HELPER to know ahout a *et of unexceptional, actual events..*

The HELPEE sets up a context by describing a world where everything, so far, is going well. Since the situation (involving HELPEE attempting to do the TASK) is presumed to be known by the HELPER, it is further assumed that the HELPER'S expectations for subsequent activity will closely parallel those of the HELPEE.

2. HELPEE wants HELPER to know about:

I) a net of execptional event which occurred*

Da net of expected, unexceptional events which did not occur.

This pattern of conducting a Helping-game is sufficiently well ingrained in the participants, that the HELPEE almost never needs to actually ask a question at this point. By simply exhibiting a failure of expectation, the HELPEE has communicated that this acts as a block to his successfully pursuing the TASK. Furthermore, he expects the HELPER to explain why this failure occurred and how he can avoid it or otherwise continue to perform the TASK.

I he third component specifies the "Treatment" phase where the HELPER communicates an explanation for the perceived failure.

'i. HELPER want HELPEE to know ahout an action which will avoid the undesired event or ca.use the desired one.*

The context description has enabled the HELPER to identify a collection of activities which he understands, and in which the HELPEE is attempting to participate. The violation-of-expectation description points out just where the HELPEE's image of the activities differs from the HELPER'S (presumably correct) image. It is from this discovered area of difference that the HELPER selects an action for the HELPEE which is expected to solve his problem.

USE OF DIALOCUE-CAMES IN THE COMPREHENSION OF DIALOGUE

In this section we describe the five stages of dialogue assimilation: nomination, recognition, instantiation, conduct, and termination, and detail the involvement of Dialogue-games with each stage.

Nomination

When dialogue participants propose a new type of interaction, they do not consistently use any single word or phrase to name the interaction. Thus we cannot determine

which Dialogue-games are involved in a particular dialogue through a simple invocation, by name (or any other foreknown collection of words or phrases). Instead the dialogue type is communicated by attempts to establish various entities as the values of the Parameters of the desired Dialogue-game. Thus, an utterance which is comprehended as associating an entity (a person or a concept) with a Parameter of a Dialogue-game suggests that Dialogue-game as a possibility for initiation. Nomination can take place through either the Match and Deduce processes or the Proteus process described below, depending on details of the input.

Recognition

The nomination processes are reasonably unselective and may activate a number of possible Dialogue-games, some of which may be mutually incompatible or otherwise inappropriate. There is a process called the Dialogue-game Manager, which investigates each of the nominated Dialogue-games, verifying inferences based of the Parameter Specifications, and eliminating those Dialogue-games for which one or more Specifications are contradicted. This approach sidesteps the issues surrounding backing-up and retrying. All hypotheses are "carried forward" until they are no longer plausible.

I nntnntiation

Once a proposed Dialogue-game has successfully survived the filtering processes described above, it is then instantiated by the Dialogue-game Manager. Those Parameter Specifications not previously known (represented in the Workspace) are established as new inferred knowledge about the Parameters. It is through these instantiation processes that a large part of the implicit communication between participants of the dialogue is modeled.

To illustrate this, suppose that the following are represented in Workspace (i.e., known):

SPEAKER does not know how to do a TASK.
SPEAKER wants to know how to do that TASK.
SPEAKER wants to do the TASK.

These are adequate to nominate the Helping-game. In the process of instantiating this Dialogue-game, the following predicates are added to Workspace:

SPEAKER believes HEARER knows how to do TASK.
SPEAKER believes HEARER is able to tell him how to do TASK.
SPEAKER believes HEARER is willing to tell him how to do TASK.
SPEAKER wants HEARER to tell him how to do TASK.
SPEAKER expects HEARER to tell him how to do TASK.

The model, then, predicts that these predicates are implicitly communicated by an utterance which succeeds in instantiating the Helping-game. This corresponds to a dialogue in which "I can't get this thing to work" is taken to be a request for help, even though on the surface it is an assertion about ability.

Conduct

Once a Dialogue-game is instantiated, the Dialogue-games Manager is guided by the Components, in comprehending the rest of the dialogue. These Components are goals for the dialogue participants. For the speaker, these goals guide what he is next to say; for the hearer, these provide expectations for the functions to be served by the speaker's subsequent utterances.

These "tactical" goals *are* central to our theory of language: an utterance is not deemed to be comprehended until some direct consequence of it is seen as serving a goal imputed to the speaker. Furthermore, although the goals of the Components are active only within the conduct of a particular game, they are so constituted that their pursuit satisfies the goals described in the Parameter Specifications, which were held by the participants prior to the evocation of the Dialogue-game.

In the case of the Helping-game, the goals in the "diagnostic" phase are that the HELPEE describe a sequence of related, unexceptional events leading up to a failure of his expectations. These goals are part of the state of the HELPER as he assimilates this initial part of the dialogue, both in that he knows how the HELPEE is attempting to describe his problem, and also that the HELPER knows when this phase is past, and the time has come (the "treatment" phase) for him to provide the help which has been implicitly requested.

Termination

The processes described above model the identification and pursuit of Dialogue-games. How, then, are they terminated? As indicated previously, the Parameter Specifications represent those aspects of dialogues that are constant over that particular type of dialogue. The Dialogue-game Model pushes this a step further in representing that the dialogue type continues *only aS long at* the Parameter Specifications continue to hold. Whenever any Specification predicate ceases to hold, then the Model predicts the impending termination of this Dialogue-game.

For example, if the HELPEE no longer wants to perform the TASK (either by accomplishing it or by abandoning that goal), then the Helping Dialogue-game terminates, corresponding to the concurrent termination of the helping interaction. If the HELPER becomes unwilling to give help, or discovers that he is unable, then the Helping-game also terminates. Again, we have one simple rule that covers a diversity of cases--a rule for termination that captures the variety of ways that the dialogues we have studied end.

A *HELPING DIALOGUE*

Let us briefly consider a naturally occurring helping dialogue to illustrate the operation of the Dialogue Model.

1.1 L SORRY TO BOTHER YOU AGAIN...
1.2 I An HAVING SOME TROUBLE UITH RUNOFF
1.3 I AM TRYING TO GET A COPY OF THE PROCESSED
1.4 FILE BACK IN MY DIRECTORY, BUT IT SEEMS TO
1.5 BE DIVERTED TO THE PRINTER.
1.6 WHAT CAN I DO? GO AHEAD.

2.1 0 JUST A MIN..
[operator executes the runoff procedure]

2.2 YOU MUST TYPE A COMMA AFTER INPUT FILE NAME

2.3 IN ORDER TO SPECIFY AN OUTPUT FILE

3.1 L THANKS.. IS THERE ANY WAY TO GET THE PRINTED
3.2 OUTPUT BACK OR AT LEAST THROUN AWAY?
3.3 THERE wAS FAIRLY IMPORTANT AND SOMEWHAT

3.4 CONFIDENTIAL INFO IN THERE...GO AHEAD

4.1 0 UE SHALL SHRED THE OUTPUT..OK

5.1 L THANKS... BYE

6.1 0 BR;BYE

Figure 3. A helping dialogue

The Helping DG is bid in line 1.2, with the linker L asserting inability to use the program Runoff satisfactorily. L then launches into the pursuit of the Helping DG in lines 1.3 through 1.6, stating a context for his problems in 1.3 and a violation of his expectation in 1.4 and 1.5. Since this completes the "diagnosis" stage of Helping, he ends his turn. Since the operator 0 can at that point identify the problem (after some preliminary information gathering effort), he can perform the "treatment" stage of Helping in lines 2.2 and 2.3, instructing the linker to take certain actions to avoid the problem. The Helping DG now terminates, as indicated by the "thanks" on line 3.1, since the immediate goal of the linker has been satisfied.

A second Dialogue-game, Action-seeking, is initiated following this on line 3.1. Just like the initiation of Helping on line 1.2, this bid for initiation is indirect, since the question on lines 3.1 and 3.2 is on the surface only simple information seeking. However, since this question establishes the parameters of the Action-seeking Dialogue-game, it is comprehended as a bid for that DG, which is accepted and pursued by 0 in line 4.1. This Dialogue-game terminates in line 5.1, since the immediate goal motivating the Action-seeking has been achieved. L bids a termination of the interaction ("bye"), indicating that he has no more goals he feels he can further through dialogue with 0, and this bid is accepted by 0 in 6.1.

SYSTEM DESIGN

The Dialogue Model System represents a participant by two memories and six principal processes, as shown in Figure 4. The long term memory (LTM) is the system's model of a participant's knowledge prior to the dialogue, including both knowledge of language and knowledge of the world. The Workspace is the model of his current awareness, including actual elements of the dialogue, various analysis products of them, his model of his own goals and knowledge and his model of his partner's goals and knowledge. Both memories use the same semantic network representation, which in its formal manifestation consists primarily of nested predicates.

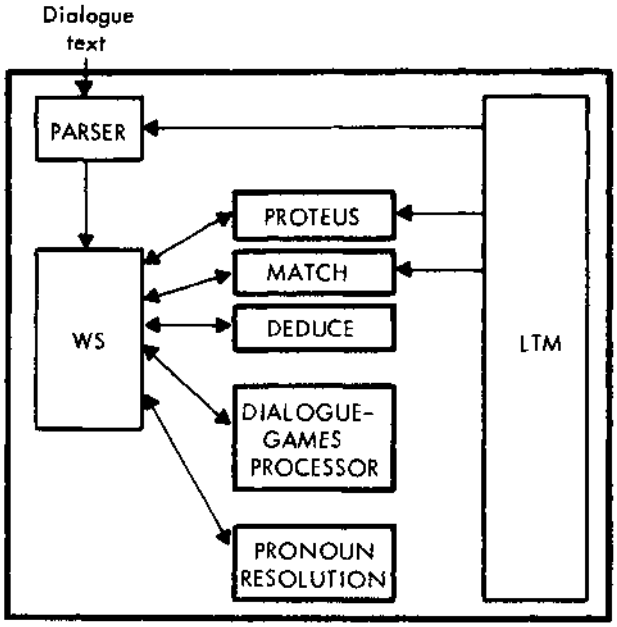

Figure 4. Dialogue Model System

All of the processes deposit their results in the Workspace, and all but the Parser take their inputs from the Workspace.

The Parser represents each successive utterance from the surface string using a standard ATN *Grammar similar* to those developed by Woods (1970) and Norman, Rumelhart, & the LNR Research Group (1975). We use a case grammar representation, with each utterance specified as a main predicate with a set of Parameters. Because this module is a conventional parser whose implementation is well understood, we have so far produced hand parses of the input utterances, following an ATN grammar.

The pair of processes called Match and Deduce function primarily as a production system. Match identifies rules in LTM whose Condition parts newly match parts of the Workspace, and Deduce applying these rules to change the Workspace according to their Action parts.

Proteus is a spreading-activation process which nominates knowledge structures of LTM for attention in the Workspace based on the connectivity and specific adjacencies of LTM objects to objects found in both LTM and the Workspace.

There is a collection of Pronoun-processes, such as the You-process and the It-process, that resolve pronouns found in the input by creating new Workspace objects or merging existing ones.

The most complex of the processes is the Dialogue-game Manager, which controls awareness of the current set of active Dialogue-games. It identifies those input turns which contain bids of games and acceptances of games, creates instances of the games in the Workspace, asserts in the Workspace the implicit communications which arise from initiation of a game by the partner, activates game Components, identifies games whose termination conditions are satisfied, identifies bids and acceptances of terminations, and closes terminated games.

The most detailed exposition of the operation of the Dialogue Modeling System is to be found in an extensive hand-simulation of the system in Levin & Moore 1977).

COMMUNICATION PHENOMENA AND MODEL OPERATIONS

How does this model expand the set of language phenomena that are explainable by process models? Many correspondences can be found between model operations and language phenomena, at a wide range of levels of detail. We suggest the correspondences for three:

1. topic structure

2. the sense of having comprehended

3. implicit communication

We then indicate the model's significance in man-machine communication design.

For each participant, the model identifies particular dialogue games in which he is actively engaged at particular points in the dialogue. The games of the moment are nested, and each has a parameter called "Topic." Topics come and go with their games. The degree of correspondence between peoples' perceptions of topic change and the entry and exit of these Topics is an empirical issue. Our informal judgment is that they correspond well for a high proportion of cases, and we plan to perform formal tests to establish this.

The theory (as represented by the model) says *in* effect that topic is not an independent attribute of dialogue, but rather is tightly coupled to simultaneously occurring participants' goals and the conventions by which such goals combine. For machine comprehension, this means that topic can be recognized more effectively by seeking the joint configuration of goal and topic rather than by independent topic assessment methods such as Keyword-based topic identification schemes. It also suggests why the keyword schemes are as limited and unreliable as they are. Establishing this close coupling between topic and goals would have significant consequences in discourse theory, since it gives a basis for discerning the gross structure of one of the most irregular genres of text: transcripts of naturally occurring dialogue.

In the model, and in people's accounts of their attempts at comprehending text, some text is seen as comprehensible and other text is not. Part of the analysis of each utterance is to regard it as an act, to identify the particular kind of act being performed, and to identify the known goals which it serves. If the act serves a goal which (in the view of the hearer) is held by the speaker of the utterance, then the model asserts a proposition in the workspace using the "Comprehends" predicate. The "Comprehends" predicate takes three arguments: an utterance, a person who spoke the utterance, and a goal.

When this proposition is asserted, the analysis of the current utterance stops and the next utterance is processed by the Parser. This "stopping rule" is a selective method for limiting the amount of inference applied to utterances. Since deciding when to stop has been a troublesome problem in other language processors, we feel that it is a significant contribution. (Schank & Rieger, 1974)

The assertion in the hearer's workspace that he "Comprehends" is the internal event which corresponds to a person's sense of having comprehended the utterance. The model fails to assert a "Comprehends" proposition not only for nonsense, but for semantically well formed non sequiturs as well. Where there are alternative semantically consistent interpretations, the requirement that the text be motivated provides a new basis for selecting correct interpretations.

Of course, the adequacy of this view of the sense of comprehension is an empirical issue subject to test.

Implicit communication is a pervasive feature of natural language use, but is by its nature difficult to reach with syntactic and semantic analyses. Ordinarily a person who asks a question as a turn in a dialogue conveys not only the question, but also several assertions, including 1) that the speaker wants the hearer to answer the question, and 2) that he regards the hearer as able to answer it. (In contrast, the same question might be asked rhetorically in another context without conveying either of these assertions.) Indirect speech acts also perform implicit communication, as, for example, questions posed as assertions:

(1) "I'd like to know why you're late."

The model derives the implicitly communicated content in two ways: by asserting the parameter specifications whenever a game is entered, and by developing motivational explanations for utterances. The first was exemplified in the discussion of the HELPING-GAME. For the latter, consider the difference between

(2) "I'm taking a survey. Do you have a match[9]"

and

(3) "I've lost my lighter. Do you have a match?"

In some rather ordinary contexts, the occurrence of (2) would be seen as an information request, in which "Do you have a match?" functions as a simple question. In contrast, (3) would be seen as a request. A simple reply of "Yes" would be seen as cooperative for (2) but for (3) it would be seen as uncooperative because it was unresponsive to the implicit request.

We would expect the model to correctly respond in these two different ways. For (2), it would posit the goal of knowing whether the hearer has a match, and regard it is a plausible goal of a survey taker. The question would be assimilated as pursuit of that goal, and not seen as indirect. For (3) there is no corresponding goal, and so the question would be identified with a precondition of an action (giving a match) and thereby be analyzed as indirect. This result is preferable to any sort of analysis that treats the questions in (2) and (3) alike. The example also illustrates how inference limitation is often necessary to correct assimilation of implicit communication.

There is a striking contrast between this model of communication and what ordinarily happens in man-machine interaction. Universally, the man-machine interface does not analyze its input to identify the goals of its user. There are no conventions of mutual goal-pursuit, no discussion of the means-ends relationships of the interface. All of the important issues of how to do things and what to do next must

be resolved in the mind of the user, because the interface simply has nothing to do with goals.

These interfaces impress many people as alien and difficult to communicate with, and they are a formidable source of difficulty for the computer-naive. (Mann, 1975) Since they lack the goal-based superstructure of human communication, the user is blocked from using a substantial and central portion of his communication skills, and may thereby be blocked from reaching his goals at all.

The lack of this superstructure seems to be a deep-seated source of difficulty, one that could with some care be corrected by extensions of the formal languages currently in use. Certainly an interface that knew what the user was trying to accomplish could be useful in many ways that are presently infeasible.

/I *PERFORMANCE BENCHMARK FOR THE MODEL*

The performance of this model can be assessed in an unlimited number of ways in response to specific interests. Without some specific notion of what kinds of performance are of interest, we would have no basis for claiming that the model is correct, effective, working, or better than previous models. The common practise of simply presenting examples does not yield a significant performance assessment, in part because it fails to characterize the range of inputs over which the model is effective.

We have attempted to develop a more suitable basis for performance assessment which would identify phenomena of interest explicitly and operationally. This part of the research was actually done before the model design described previously and our modeling was guided by the knowledge gained in learning to make reliable judgments of the phenomena.

One of the tasks of the Dialogue Model System is to follow the changes of topic in a dialogue. Since topic is not indicated formally in English-language conversation, human judgments of how topics occur must be used. Success in modeling topic phenomena is indicated by *correspondence* between model behavior and human judgment s on the tame text.*

Some kinds of human judgments about text *are* much more widely shared than others. For assessing models, we are particularly interested in comparing the model performance with judgments which are reliable (produced by a large proportion of those who judge the text), reproducible (by other researchers), and independent of the knowledge domain of the text.

To obtain these judgments we have developed a Topic Observation Method which people can use to annotate dialogue text (Mann, et al 1975). The method gives systematic instructions for annotating topic beginnings and topic endings for each of the individual participants. Observers show the scopes of topics and assign labels which indicate when topics are resumed or shared by the two parties. The annotations thus specify topological attributes of the text. Figure 5 shows *an* example of topic annotation of a dialogue.

In the course of developing this method, we discovered several interesting properties of people's knowledge of Topic. People can identify the topics relative to an individual

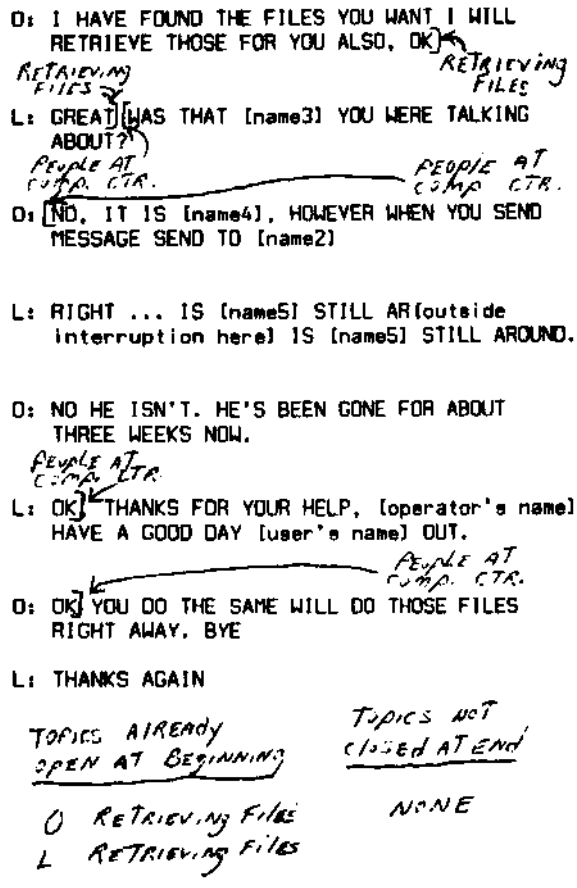

Figure 5. Dialogue annotated for topic

Figure 5. Dialogue annotated for topic participant, but we found no reliable way to have them assign the "current topic" of the dialogue as a whole. Also, there is a significant fraction of dialogue devoted to negotiating what the following topic shall be. People label topics idiosyncratically, so that what one person regards as a topic of "Medical Ethics" another person regards as a topic of "Doctor-Owned Hospitals." On the other hand, judgment of where topic, changes occur is not idiosyncratic. This method was developed as part of a group of annotation methods intended to operationalize particular human communication phenomena (Mann, Moore, Levin & Carlisle, 1975). The others are:

- Repeated Reference (multiple mentions of things)
- Requests (including initial requests, responses and dispositions)
- Similar Expressions (in-context paraphrase-like functional equivalence)
- Correction Events
- Expressions of Comprehension

For each of these, including Topic Structure, a formal test of inter-observer reliability was performed (Mann, Carlisle, Moore & Levin, 1977). We found that the method could be applied with very high levels of agreement among different Observers. In the course of developing these six categories, about fifteen other proposed category definitions were rejected, primarily because of unreliability.

In the Model, Topic changes are seen as consequences of changes in the set of active Dialogue-games headed by individual participants. (This incidentally explains why topic judgments must be made relative to individual participants.) For the model to explain the topic structure of a dialogue, the changes in the games in the Workspace of the hearer should correspond to the Observer's annotations of topic beginnings and endings. For significance, the model actions should correspond to annotations for a suitably large body of text.

So far the model has been tested only informally, but an extensive formal test is planned. Whatever the outcome, the Topic Structure Observation Method exhibits a class of human judgments about text which a theory of natural language communication, whether composed of processes or not, should explain.

Each of these Observation Methods provides a way to factor the complex problem of understanding natural language into simpler problems of reproducing and accounting for particular language-related phenomena. Such methods of factoring and objectifying the problem are sorely needed in this complex technical domain, and in artificial intelligence in general.

RKFKRKNCES

Air.tin, J. L. *How to do things with words.* Cambridge, MA: Harvard University Press, 1962.

Charmak, E. Organisation and inference in a frame-like system of common sense knowledge. In R, Schank & B. L Nash-Webber (Eds.), *Theoretical issues in natural language processing.* Cambridge, MA: Bolt, Beranek and Newman, Inc., 1975.

Deutscb, *B.* G The structure of task-oriented dialogues. In L. D. Lrman (Ld.), *Proceedings of the ILEE Symposium on Speech Recognition.* Pittsburgh, PA: Carnegie-Mellon Univ., 1974.

Cordon, P., & I akoff, G Conversational postulates. *Papers from the Seventh Regional Meeting,* Chicago Linguistic Society, 1971.

Grice, M. P. Logic, and conversation. In P. Cole & J. L. Morgan (Ids.), *Syntax and semantics.* New York: Academic Press, 1975.

Grosz, B. J. Discourse Analysis. In D. E. Walker (Ed.), *Speech Understanding Research.* Menlo Park, CA: Stanford Research Institute, 1976.

Gros/, (3, J. The representation and use of focus in dialog understanding. Doctoral dissertation (forthcoming), Berkeley, CA: Univ. of Calif., 1977.

Labov, W., & Fanshel, D. *Therapeutic discourse: Psychotherapy as conversation.* Draft copy, 1974.

Levin, J. A. Proteus: An activation framework for cognitive process models. Unpublished doctoral dissertation, San Diego, CA: Univ. of Calif, San Diego, 1976.

Levin, J. A., & Moore, J. A. Dialogue-games: Meta-comrnunication structures for natural language interaction (I5I/RR-77-53). Marina del Rey, CA: Information Sciences Institute, 1977, In press, *Cognitive Science,* 1977, i. 4.

Mann, W. C. Man-machine communication research: Final report (ISI/RR-77-57). Manna del Rey, CA: Information Sciences Institute, 1977.

Mann, W C. Why things *arv* so bad for the computer-naive user (ISI/RR-75-32). Marina del Rey, CA: Information Sciences Institute, 1975.

Mann, W. C, Carlisle, J. K, Moore, J. A., & Levin, J. A. An assessment of reliability of dialogue annotation instructions (ISI/RR-77-54). Marina del Rey, CA: Information Sciences Institute, 1977.

Mann, W. C, Moore, J. A., Levin, J. A., & Carlisle, J. H. Observation methods lor human dialogue (ISI/RR-75-33). Marina del Rey, CA: Information Sciences Institute, 1975.

Mirv.ky, M. A framework for representing knowledge. In P. H. Winston (Ed), *The psychology of computer vision.* New York: McGraw-Hill, 1975.

Moore, J. A., Levin, J. A., Mann, W. C, A Goal-Oriented Model of Human Dialogue, submitted to *American Journal of Computational Linguistics,* May 1977.

Newell, A., & Simon, H. A. *Human problem solving.* Englcv/ood Cliffe., NJ: Prentice-Hall, 1972.

Rieger, C. The commonsense algorithm as a basis for computer models of human memory, inference, belief and contextual language comprehension. In R. Schank & H. I. Nash-Webber (Eds.), *Theoretical issues in natural language processing.* Cambridge, MA: Bolt, Beranek and Newman, Inc., 19/5.

Rumrlhart, D. E. Notes on a schema for stories. In D. G. Bobrow & A. Collins (Eds.), *Representation and understanding: Studies in cognitive science.* New York: Academic Press, 1975.

Sacks, K, Schegloff, E. A., & Jefferson, G. A simplest systematics for the organization of turn-taking for conversation. *Language,* 1974,50,696-735.

Schank, R. C, K Abehon, R. P. Scripts, plans and knowledge. Paper presented at the Fourth International Joint Conference on Artificial Intelligence, Tbilisi, USSR, August 1975.

Schank, R C. & Rieger, C. J. Inference and the computer understanding of natural language. *Artificial Intelligence,* 1974, 5, 373-412.

Searle, J. R. *Speech acts: An essay in the philosophy of language.* Cambridge, England: Cambridge University Press, 1969.

Searle, J. R. Indirect speech acts. In P. Cole K J. L. Morgan (Eds.), *Syntar and semantics.* New York: Academic Press, 1975.

Shatz, M. How young children respond to language: Procedures for answering. *Papers and Reports on Child Language Development,* Palo Alto, CA: Stanford University, 1975.

Thorndyke, P. W. Cognitive structures in comprehension and memory of narrative discourse. *Cognitive Psychology,* 1977, Q, 77-110.

Wittgenstein, L. *Philosophical investigations* (3rd ed.). New York: Macmillan, 1958.

Woods, W. A. Transition network grammars for natural language analysis. *Communications of the ACM,* 1970, *13,* 59^-606.