

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
ARTIFICIAL INTELLIGENCE LABORATORY

AI Memo No. 413

May 1977

Levels of Complexity in Discourse

For Reference Disambiguation and Speech Act Interpretation

by

Candace Bullwinkle

Abstract: This paper presents a discussion of means of describing the discourse and its components which makes speech act interpretation and reference disambiguation possible with minimal search of the knowledge in the database. A portion of this paper will consider how a frames representation of sentences and common sense knowledge provides a mechanism for representing the postulated discourse components. Finally some discussion of the use of the discourse model and of frames in a discourse understanding program for a personal assistant will be presented.

This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the laboratory's artificial intelligence research is provided in part by the Advanced Research Projects Agency of the Department of Defense under the Office of Naval Research under Contract Number N00014-75-C-0643.

Introduction The person who communicates with a personal assistant, whether human or machine, wants to request some action of the assistant via sentences in English. Generally, a single sentence is insufficient to capture all the information that is to be given as well as an unnatural way to make a request. However, as several example dialogues below will show, the human user does not tightly relate the sentences s/he speaks about a particular subject. It is instead the job of the hearer to interpret how the incoming sentence is related to the previous discourse.

Each sentence or clause of a discourse that makes some demand upon the hearer must be interpreted for the kind of demand being made. These demands are generally referred to as speech acts {1}. However, the speech acts are not just strings of individual requests. They have a connecting pattern which the hearer must extract as the discourse goes on. Thus the hearer's task is two-fold: to interpret the speech act in a clause and to relate that speech act to the overall discourse. In this paper this two-pronged task will be referred to as speech act interpretation.

Closely associated with speech act interpretation is the process of understanding what the various noun and pronoun phrases of the incoming sentence refer to. Speakers denote previously mentioned objects in a variety of ways with apparent ambiguity in the choice of referents. Objects associated with some previously mentioned object (such as *the time when a meeting* has been mentioned) are not marked in any linguistic manner, and the hearer must decide what if any connection exists between the two. This process, which will be called reference disambiguation, creates pieces of information about the relation between terms in the current sentence and

those in previous ones.

The two processes described above are central mechanisms in discourse interpretation. In this paper I will discuss a means of describing the discourse and its components which makes speech act interpretation and reference disambiguation possible with minimal search of the knowledge in the database. I will show how speech act interpretation plays a vital role in reference disambiguation. In a portion of this paper, I will discuss how a frames representation of sentences and of common sense knowledge provides a representation mechanism for the postulated discourse components. Finally some discussion of the use of the discourse model and frames representation in a discourse understanding program for a personal assistant will be presented.

The two processes of speech act interpretation and reference disambiguation require a detailed descriptive representation as the structure of the discourses becomes more complex. By looking at several sample discourses of increasing complexity, more can be seen about the kinds of capabilities which must be included in reference disambiguation and speech act interpretation. Figures 1 and 2 present various problems in reference disambiguation and speech act interpretation which will be discussed by examples. The problems are ordered by increasing complexity, and this complexity will be considered in the section following the figures. Special terms like *topic*, *surface inference*, *contrastive reference* and *stacked topic* will be defined in the course of this paper.

Reference

Problem	Example
-choosing the referent of a unique proper name	1
-reference of definite noun phrases where the noun phrase is previously used	1
-anaphoric references of the topic of a request	2
-reference of definite noun phrases where the reference is to a concept associated with the topic via surface inferences	2
-anaphoric references of objects other than the topic	2
-reference of definite noun phrase described in previous discourses	5
-reference of definite noun phrases where the discourse implies the concept to which the noun phrase refers	4a
-contrastive references	6
-anaphoric references that refer to unmentioned objects	7
-references to stacked topics	
-anaphoric references that refer to unlikely topics	
-reference of definite noun phrases where the reference is to a concept associated with the topic via deep inferences	

Figure 1

Speech Acts

Problem	Example
-recognition of a speech act when no previous discourse exists	1
-recognition of a speech act which is related to the discourse by the same topic	1
-recognition of a speech act which is related to the discourse but which appears to have a different topic	3a, 3b
-recognition of the role of assertions and questions to other speech acts	4a, 8
-recognition of multiple speech acts of the same type with different topics	6, 7
-recognition of speech acts that appear unrelated to the discourse and have different topics but which are related by deep inferencing	

Figure 2

Some Discourse Examples In the examples below, different aspects of the problem of determining both the referent of definite noun phrases and anaphoric phrases will be presented. These examples will also show what kinds of interactions can occur between different speech acts and how these interactions are reflected in the references.

In (1) below, there are two kinds of speech acts. The first is a request for some action to be performed, i.e. scheduling a meeting. The sentences that follow contain a second kind of speech act, an assertion about when the meeting ought to occur. A speech act has a central concept which the speech act is about. In this paper this central concept is called the topic; here the topic is the meeting. A basic problem of speech act interpretation is to determine what topic a particular speech act refers to. Once a topic can be chosen for a speech act, the topic becomes the focal point of reference resolution. Topic choice is difficult because two different speech acts may be related only by reference to the same topic, and yet the topic may not be explicitly stated in both. The references in this example range from easily resolvable ones like *I* and *a meeting*, to simple anaphora (*it*) and to references requiring some database search (*Bruce, Bruce's office*). Later a frames representation will be presented to permit reference and anaphor resolution with a minimum of, or no, search.

- (1) I want to schedule a meeting with Ira. It should be at 3 p.m. tomorrow. We can meet in Bruce's office.

In (2), the speech acts are the same as in (1), but an object associated with the meeting, the time, is not clearly stated as referring to the time of the meeting. Instead this is assumed to be determinable

from the context and the hearer's knowledge of meetings. Also in (2), the referent of *his* ought to be resolvable with a minimum of search.

(2) I want to schedule a meeting with Prof. Goldstein. The time should be 4 p.m. on Thursday. We can meet in his office.

In (3a) and (3b), a new speech act, *invite*, has been introduced. This speech act makes a request, and there is no clear indication of what the inviting has to do with the previous context. Again it is left to the hearer to determine from what s/he knows about meetings and scheduling a meeting to decide what the connection is.

(3a) Ira wants to have a party at his house. It's going to be at 8 o'clock p.m. Invite {everyone at the lab; Mitch, Dave, Candy, Bruce and Beth} to attend.

(3b) Ira wants to have a party at his house. It's at 25 Wildwood St. Invite {everyone at the lab; Mitch, Dave, Candy, Bruce and Beth} to attend.

(3a) and (3b) add further complexity to the use of reference disambiguation. The *It's* in sentence 2 of (3a) and (3b) refers to *party*. Most existing systems misinterpret the anaphoric reference in (3b) because *house* occurs between the pronoun and its referent. Correct disambiguation requires some knowledge of what has been talked about so far, which is the topic, and how sentence 2 relates to this topic. Sentence 3 of (3a) and (3b) shows two alternative ways to tell who is to be invited. One is a string of names which can be disambiguated just as all the previous examples of names are. The other uses a quantifier over a set. While the current work does not claim a solution to the quantifier problem, it is introduced here to show what complexity can be expressed in the most simple looking discourses. Finally, the ellipsis of what is being attended by everyone must be resolved to determine what relation the people have to the

invitation.

Example (4a) points out two new problems. The first 3 sentences, though similar to those above, offer a new reference problem. The first sentence implies a meeting to which the phrase *our meeting* in sentence 2 refers. By the use of topic, this relationship can be determined without searching the database of all meetings for the referent. The last sentence of (4a) points out another feature of speech acts. It requests the answer to a question which is not independent of the context. In general, questions may or may not take this form. An acceptable alternative discourse which does not have this form would be one like (4b).

(4a) I want to get together with Bruce. Our meeting should be before Friday. We can meet in my office. Can Mitch attend also?

(4b) I want to get together with Bruce. Our meeting should be before Friday. We can meet in my office. Can I also attend the staff meeting this week?

In example (5) a new speech act is used, but this speech act is not significantly different from the speech act of *schedule*. This example illustrates the reference problem of *the pa-meeting*. While no previous discourse exists in which to find the referent of this term, the term may have a unique referent even when several *pa-meetings* have been scheduled (due to the participants, places and times). This suggests that previous discourses must be saved in a form suitable for providing necessary information to disambiguate the reference.

(5) I want to re-schedule the *pa-meeting*. Change it to Thursday.

The next several discourses present examples of speech acts which have different topics and different discourse structure, and yet the speech acts being requested are the same. The first of these, (6), has a structure

with two different parallel requests about the same kind of events. It uses a kind of contrastive referencing, *the one...the other*, that is often used for showing contrast between two similar notions. In addition there is an assertion about the first request after the statement of the second request. The second example, (7), also requests two separate scheduling events, but the requests have a sequential form. These discourses suggest that there may be several ways to structure speech act requests to convey the request completely but with minimal repetition of information. Consideration of discourse structure when interpreting these requests assists in reference disambiguation: the structure shows whether one or more objects are being discussed within the context.

- (6) I want to schedule two meetings. One should be on Thursday at 3 with Ira. The other is with Bill on the 14th at 4. The Thursday one may last quite a while so schedule 2 hours for it.
- (7) Next week on Monday, I want to meet with Bruce. Please arrange it for us. Also you scheduled a meeting on Monday at 8 a.m. That's too early. Make it 10 a.m. at the earliest.

In addition to the speech act structure, example (7) points to an interesting problem. There is no previous referent for the use of *it* in sentence 2 of (7){2}. This, as well as other examples not reproduced here, suggest that special techniques are needed for finding anaphoric referents that fail to refer to some syntactic object. Such anaphors are generally understood. Inferencing about the topic of the context as a nominal similar to those in (4a) will permit the correct association.

Example (8) contains two speech acts, one a fact about some work being done and the other a request for a meeting time. The two speech acts are related because the first is a reason for the second. A similar relation

exists between the clauses of sentence 4 of example (6). As with (6) and (7), determination of the possible structures of discourse aids in the interpretation of the speech acts of (8).

- (8) Dave, Bruce and I are working on a new program together. We need to meet once a week for a month. Schedule a regular meeting time for us.

In the last example below, the phrase *the last meeting* bears an important relation to the phrase *the meeting*. In order for a reference finder to determine what *the last meeting* refers to, it is necessary to know that a) the meetings are related along some linear dimension to which *last* can apply, and b) that *the last meeting* was a meeting of the PA group. This example suggests that semantic knowledge and some pragmatic assumptions must be available to a reference disambiguation program.

- (9) The PA group will want to meet next week. The meeting should be on Wednesday. The last meeting, which was at 5, was too late, so schedule this one earlier.

Sentence 3 shows another difficulty: the focus of the discourse shifts to *the last meeting* and then back to meeting via the use of *this one*. The use of *this* and *that* to point to different foci of a discourse appears to be similar to the use of *the one...the other*. Example (9), like the previous three examples, contains the use of an assertion serving as a reason for some other request. In (9), the link between assertions and requests is needed to resolve the ellipsis associated with *earlier*.

Components of Discourse Now let us turn our attention to describing components of a discourse. Discourse organization centers around the theme and the topic. The theme of the discourse is the speech act request or assertion which is being made by the discourse. It is not necessarily the

verb in a sentence; it is the request or assertion being made using a particular verb. Discourse themes can be difficult to identify, but generally in the kinds of discourses discussed here, the leading sentence suggests the discourse theme. To illustrate, note that in example (1), the theme is a scheduling request since the discourse is taken to be a request for scheduling. It is the use of a *want* modality together with *schedule* which suggests this theme. The combination of the two are necessary, for in example (10) below, a variation of (1), the use of *schedule* in the second sentence suggests an assertion, not a request for scheduling.

(10) I have a meeting with Mitch. It is scheduled for 10 a.m. on Monday.

Each sentence of a discourse has its own speech act. These must be determined from sentential verbs, so that the overall discourse speech act may be recognized. In (1) as we have seen, the first speech act is a scheduling request. This becomes the discourse theme. The second sentence has an assertion as its speech act. To relate the two speech acts and sentences, one must recognize that the assertion gives additional information about the event for which scheduling is requested. In general, any discourse theme can have associated to it certain speech acts such as assertions, questions or further requests. The speech acts can be assumed to be related because they are about the same requested event. In (1) that event is a meeting.

Topic Every speech act is assumed to be about something, that is, a speech act has a central concept it states. This concept is the topic {3}. A request for scheduling makes the request about a meeting. The assertion of information is an assertion about some object or event. The topic is not

necessarily assumed to be the neutral case of the sentence case frame (or direct object of the sentence). Often the neutral case may contain the topic. However, this is an artifact of the kinds of sentences that are considered in this article. Most of them are requests by a speaker for some action to be performed. Since the speaker as agent is well-identified (the hearer knows, in some sense, who the speaker is), the neutral or temporal or locative are candidates for the topic, the neutral being preferred over the other two. In discourses, notably stories, the agent may contain the topic. However, even in such discourses, the teller of the tale may include him/herself as agent, and then the topic will be the neutral or some other non-agent case.

When a theme of discourse is chosen, the topic of its related sentence becomes the discourse topic. Once a topic is chosen for a discourse, subsequent sentences provide more information about the topic. They may use any of the case frame slots associated with their own sentence predicates. Very often, as in the examples above, a reference to the topic will appear in the agent slot. This is consistent with the old/new distinction [Moore, 1967] because a reference to the topic, which is old information after one sentence, often appears at the beginning of a succeeding sentence.

Sub-topics A discourse may expand various aspects of a topic and then end the discussion of that topic. Alternatively, it may mention one aspect of the topic and describe it in detail. For example, in a discourse about meetings, we may want to spend several sentences specifying the time for the meeting, why that time is best and so on.

At the point in the discourse where time is being discussed instead of the topic meeting, one would like to declare a change of topic so that the assertions or requests in the subsequent sentences will be taken to be about time. However, the meeting should not be "forgotten" entirely because the discourse may eventually return to it. In this case, the meeting topic can be said to have a background role to the new topic of time. To maintain both topics, one in the foreground, and the other in the background, the topic of meeting can be thought of as being stacked so that it can be retrieved for later use (4). This process of switching from a discourse topic to a concept that is related to the topic will be called sub-topic shift, and the topic which is stacked is referred to as the stacked topic.

Recognizing the occurrence of sub-topics can be a difficult task. One might choose the following criterion for determining sub-topic shift: A sub-topic shift occurs whenever a sentence has as its topic a concept which is not identical to the existing discourse topic. In example (2), listed as (11) below, this criterion would cause sub-topic shift to time and then to office since these noun phrases are the topics of the individual sentences.

(11) I want to schedule a meeting with Prof. Goldstein. The time should be 4 p.m. on Thursday. We can meet in his office.

However, discourse (11) is about a meeting, and sentences 2 and 3 provide additional information about it. Shifting the discourse topic to time and then to office loses this connecting thread.

As the reader may have seen, the problem of sub-topic shifts is that they cannot be predicted; they are detectable only after they occur. Consider the first two sentences of (12a) below:

(12a) I want to schedule a meeting with George, Jim, Steve and Mike.
We can meet in my office. It's kind of small, but the meeting
won't last very long anyway.

In this discourse, after sentence 2, the discourse topic is still meeting. Sentence 2 has provided information about one aspect of the meeting. The first clause of sentence 3 changes the topic since it is an assertion about *my office*. To determine the relation of the sentence to the discourse, and to interpret the anaphor correctly, one must be aware that a sub-topic shift has occurred. The skeptical reader may begin to suspect a contradiction: in order to interpret *it's kind of small*, one must know a shift has occurred, but to do the shift one must have interpreted the sentence!

The solution to this problem is to be aware that a sub-topic shift is possible once sentence 2 of (12a) is interpreted. *Office* is a potential object of further discussion (as are any concepts associated to meetings). However, sentence 2 introduces *office* as a potential anaphoric referent while other concepts related to *meeting* do not have this property. In this way, the first clause of sentence 3 can be interpreted as referring to *my office*. Of course, the nature of the assertion on which it centers is significant. It must be possible for an assertion of smallness to be made about *my office* before a referential connection can be assumed and the sub-topic shift made. This implies the need for semantic knowledge about what can be predicated of objects in the domain. The knowledge is fairly limited, however, because the decision is strictly a yes-no question about a particular object, not a decision among competing objects. Should semantic knowledge reject such an association, as in (12b) below, the

discourse topic is immediately available as the referent of the anaphor.

(12b) I want to schedule a meeting with George, Jim, Steve and Mike.
We can meet in my office. It won't take more than 20 minutes.

However, semantic marker type checks are not sufficient to conclude that the sentence *it's kind of small* is about an office. The speech act of asserting must also fit within the context. For (12a), the distinction of speech acts rests on subtle criteria, the use of present tense with no modality. Since *meeting* has been discussed only in a modal context in (12a), a non-modal assertion about *meeting* would be unacceptable whereas the non-modal assertion is acceptable for *office*. Were the first part of sentence 3 in (12a) *it will be kind of small*, the hearer is likely to associate *meeting* with the anaphor *it*.

Potential sub-topics have a short lifetime. If a potential sub-topic does not become the topic as the result of processing the sentence following the one in which the potential is seen, it is dropped as a potential sub-topic. Thus in (12b) by the end of the third sentence, *my office* is dropped as a potential sub-topic. Hereafter if *office* is discussed, it cannot be referred to using *it* until some sentence re-introduces *office* as a potential anaphoric referent.

Naturally, a sub-topic must eventually be popped when the topic shifts back. In (12a), the reference *the meeting* is a reference to the stacked topic. Generally whenever a reference to the stacked topic occurs, the sub-topic is popped, and the stacked topic becomes the topic again.

To see how these discourse components fit together for a different kind of example, consider example (13a) from Charniak [1972]. Without discussing complete discourse types for stories, I will describe briefly

how topic and sub-topics are used here.

(13a) Today was Jack's birthday. Penny and Janet went to the store. They were going to get presents. Janet decided to get a top. "Don't do that," said Penny. "Jack has a top. He will make you take it back."

By means of sub-topics, the topic shifts from presents to a top which Janet will get with the potential topic of Jack's top. When the sentence *He will make you take it back* is encountered, the anaphoric *it* is checked for a possible reference to the top which Janet plans to get as well as to Jack's top. On simple semantic criteria such as semantic restrictions on the verb *take back*, either reference is permitted. However, when considering the speech act content of *he will make you take it back*, it is plausible to associate this to the context of Penny's reasons for not buying a top. Thus *it* will be taken to refer to the top that Janet is considering and the potential topic of Jack's top will be dropped. Note that the anaphor could have been flipped to Jack's top as in (13b):

(13b) Today was Jack's birthday. Penny and Janet went to the store. They were going to get presents. Janet decided to get a top. "Don't do that," said Penny. "Jack has a top. It's green with a red stripe. He likes it alot."

Here the semantic markers of the topic and potential topic accept the property of being green with a red stripe. However, an assertion about the property of the top that Janet will buy within the context of Penny's reasons is very peculiar and would be rejected.

Sub-topic shifting is only one way to delineate the references in a discourse. In some discourses, the topic is not shifted to a sub-topic. In such a discourse, the topic and a second object, which can be related to the topic in one of several ways, are discussed in parallel. To

distinguish between the topic and the second object, a different set of anaphors are used, generally of the class of *this-that* or *the one...the other*. The topic is referred to as *this* while the second object as *that*. Similarly for *the one...the other*, as can be seen in example (6).

Discourses for the Personal Assistant As part of the Personal Assistant project at the A.I. lab, a natural language understanding program, called PAL, is being developed to understand discourses such as those in examples (1) to (9). The remainder of this paper will describe how the notion of speech act, topic, theme and discourse are used in a frames representation, and how the frames structure can be used to model sentences of the discourse. In this discussion, the author will assume the reader has a knowledge of some of the concepts of case frame semantics. Those who wish to discover how these are currently implemented in PAL will find a brief description in Bullwinkle [1977]. The mapping of noun phrases into case frame and to deep frame format will also not be discussed here and can be found in Bullwinkle. Instead this article will concentrate on concepts and processes that can be described with frames and which make reference disambiguation and speech act interpretation possible.

The notion of frames used here refers to the frames representation language (FRL) developed by Goldstein and Roberts [1977]. In this system a frame is a generalized property list where each generalized property is called a slot. The slot has several properties, called keys, the most important of which is the value (\$val) key. When this key is filled, attached procedures are run to produce "automatic" additional results. A frame for a meeting, such as the one below, gives common sense information

about meetings for use by various database procedures. Frames like this exist in the database before discourses are introduced. They act as prototypes for constructing additional frames. The frame gives a-kind-of (AKO), default and preference information about the times, places, and participants are associated with the frame as well as some details about what to do when a slot is filled (\$if-added).

```
(MEETING (AKO ($VALUE (COMMUNICATION)))
  ;; scheduling has several default options
  (SCHEDULE ($DEFAULT (CANCEL)
    (RESCHEDULE)
    (BEGIN-LATER)
    (END-EARLIER)
    (POSTPONE)
    (ADVANCE)
    (CHANGE-PARTICIPANT)))
  (PARTICIPANT ($PREFERENCE ((EXISTS-CHAIRMAN?)))
    ($IF-ADDED ((SELECT-CHAIRMAN))))
  (INSTANCE ($VALUE (PA-MEETING)))
  ;; the marker => means "look in the frame-slot-key listed here
  ;; for the information needed," i.e. => is a pointer
  (TIME ($PREFERENCE ((=> PA-PROJECT (MEETING TIME) $PREFERENCE))
    ((=> IRA (MEETING TIME) $PREFERENCE))
    ((=> RBR (MEETING TIME) $PREFERENCE))))
  (PLACE ($DEFAULT ((=> PA-PROJECT (MEETING PLACE) $DEFAULT))
    ((=> IRA (MEETING PLACE) $DEFAULT))
    ((=> RBR (MEETING PLACE) $DEFAULT)))
    ($PREFERENCE ((=> IRA (MEETING PLACE) $PREFERENCE))))
  (POSSESSIVE ($DEFAULT ((PARTICIPANT))))
```

Frames exist for actions as well as objects. The frame for *schedule* given below contains information about how it is used in discourse (in the synonymy and discourse slots) as well as some pragmatic information which is discussed in Bullwinkle [1977].

```
(SCHEDULE (AKO ($VALUE (ACTION)))
  (PRAMATICS ($REQUIREMENT (PRAGMATICS-FOR-SCHEDULE)))
  (SYNONYMY ($VALUE (((ASSERT ACTIVITY))))))
  (DISCOURSE ($VALUE (((ASSERT THING))))))
  (TOPIC ($DEFAULT ((NEUT COM LOC))))
```

The discourse topic as it is used in PAL is an extra-sentential

device. It is used for finding referents that cross sentence boundaries. Intra-sentential referencing is needed as (S1) and (S2) below show.

- (S1) He said he was going.
 (S2) John washed himself.

Without intra-sentential referencing, the co-referentiality of *himself* to *John* and possible coreference of the two uses of *he* in (S1) would be impossible. Sentential referencing rules following the work of Ross[1967], Lasnik[1976] and Keyser & Postal[1976] are used.

As defined previously, the theme is the overall request which is invoked in the discourse. Each theme is represented as a frame which includes some information about what kinds of speech acts can be used to extend the discourse. This information, which I will call the discourse scenario, allows the frame for the theme to "see" an additional sentence as making an assertion about some slot of the topic or as introducing an additional request. A theme can have potentially many different kinds of topics, e.g. *schedule* as a theme can have *meetings*, *parties*, *lectures*, etc. as topics. The theme does not keep a list of all the slots of each one of these on its own frame. Rather it computes the topic slots from the topic itself which is also represented as a frame. The results of this computation are expectations of what the remaining sentences of the discourse are likely to discuss. Inferences made about sentences using these expectations are surface inferences since they use the slot structure of the topic to make the inference.

The frames structure limits the type and number of concepts that can be discussed for a given topic. It limits the concepts to those that are slot-values for the topic frame. Rosenberg [1977] call this limited set

the context horizon of the topic. This structural limitation has a useful effect since it provides a limited number of ways to interpret subsequent sentences. This limitation makes necessary the process of sub-topic shift so that a sentence about one of the slot-values (for example, size) of the slot-value (for example, place) of the topic (for example, meeting) can be processed. Yet while the context horizon limitation forces the additional mechanism of sub-topic shift to be used in discourse interpretation, it also focuses the sub-topic shift to a limited set of concepts, i.e. the slots of the topic. Hence, while the context horizon limitation is a limit on the structure of frames and discourse, it is a productive limitation.

Since the initial stages of the research presented here, a notion of context has been employed. At first context was simply a list of all frames that were built at any one of the language understanding modules {5}. It was believed that this data structure provided a list of what could and could not be referenced by anaphora and definite noun phrases. However, further research suggests that the topic itself provides the structure for anaphoric references, while the topic together with expectations computable from the theme and topic provide for most definite references. Thus, the topic and the theme serve as a context mechanism in and of themselves. To see why a simple list of frames is inadequate as a context, consider example (14) below.

(14)* George wants me to have a meeting with Sally this Tuesday.
It will be at 3 p.m. We can meet in my office. Before the
meeting, Bill will ask him to speak.

Most informants tell me that they do not immediately have a referent for the phrase *him* in the last sentence, and that they have to re-read the

discourse for something that could be the referent. If a simple list of frames is used, a frame for *George* will be on the list of possible referents and be chosen as referent since *Bill* will be excluded due to intra-sentential reflexivization rules. However, if a theme-topic model is used, the topic is the meeting and *George* does not appear in any of the slots of meeting, thereby making it difficult to refer to him with an anaphoric reference.

Using the topic and theme notions and frames as described above, it is possible to determine the references of noun and pronoun phrases in the examples (1) to (9) above, as well as interpret the speech acts of these examples. In finding referents a minimum of search is done, and the search is directed by knowledge of what the discourse is about. To see how this is done, examples (1) and (2) will be sketched in detail, and then a briefer sketch of the solution to some problems in the remaining examples will be given.

- (1) I want to schedule a meeting with *Ira*. It should be at 3 p.m. tomorrow. We can meet in *Bruce's* office.

Following interpretation of sentence 1 of (1), the theme is set to the frame of *schedule* and the topic to *meeting*. A referent for *I* is assumed to be the speaker. A referent for *Ira* is chosen from the instances of people via a program that knows the syntax and semantics of proper names. In sentence 2, the referent of *it* is chosen to be the topic because the topic passes gender, number and simple semantic agreement checks.

The discourse scenario for the theme *schedule* accepts assertions which provide additional information about the topic. Since the time phrase is additional information about the topic meeting, sentence 2 is judged to be

part of the schedule discourse. Likewise, sentence 3 is an assertion of information that can be associated with the topic, and this sentence is taken to assert the place of the topic. Sentence 3 exemplifies PAL's ability to use speech acts that do not directly specify the topic to which they refer.

The resolution of *Bruce's office* in (1) occurs in two parts; first a referent for *Bruce* is found via the name program mentioned above. If *Bruce's office* is not mentioned as a slot in the frame for *Bruce*, then an office is searched for among the instances of offices in the frames database such that this office has *Bruce* as an occupant or owner. The search is minimal because it occurs among a relatively small collection of likely objects rather than all the objects that exist in the database.

(2) I want to schedule a meeting with Prof. Goldstein. The time should be 4 p.m. on Thursday. We can meet in his office.

The processing for example (2) sentence 1 proceeds as in (1), sentence 1. In processing sentence 2, the reference finder expects the topic to be either the referent of *the time* or to contain a slot which *the time* refers to. This use of surface inferencing for definite noun phrases is possible because of the slot structure of frames. Processing of sentence 3 is similar to example (1) except for the reference disambiguation of *his*. The referent is found by considering instantiated slots in the topic frame which contain additional information about the topic, that is, the time and participant slots. Since time slots fail a simple gender check, a referent is found among the values listed on the participant slot.

From example (2) it can be seen that the semantics of the reference finder are somewhat limited. Ideally, a reference finder should consider

only those slots of a frame which can be filled by objects of the same semantic type as the pronoun. This is possible, however, only if there is knowledge in a frame about what kinds of objects could potentially fill a slot of the frame. This knowledge has not yet been incorporated into the frames for PAL and the reference finder.

Example (3) shows the first use of intra-sentential referencing for determining the co-reference between *his* and *Ira* by means of the precede and command rule of Lasnik[1976]. PAL chooses *the party* of sentence 1 as the referent of the *it* in sentence 2 because *party* is the topic. *The house* as a referent is overlooked entirely because it plays no part as the topic and because it is not a potential topic.

To understand the last sentence of (3a), PAL is currently being programmed to understand requests within a schedule request. This will permit PAL to interpret requests for invitations as well as re-statements of the initial discourse request, which often occur in short discourses. Using the discourse scenario in this way, PAL will be able to conclude that the invitation request is a request for the guests to attend the event denoted by the topic. This same method, when applied to questions, will create the link between the scheduling discourse and the question in example (4a).

Example (5) requires the addition of other speech act requests to PAL, using frames and discourse scenarios in the same manner as a scheduling request. It also requires, as was pointed out earlier in this paper, that previous discourse be saved in some organized fashion. This requirement is currently handled in the following way. Any topic instantiated by a

discourse is an instance of the prototype frame of which it is a kind. Thus a pa-meeting frame was generated in the discourse which occurred sometime before discourse (5). That pa-meeting, call it pa-meeting35, is an instance of the general pa-meeting frame. When a referent is sought for the pa-meeting in (5), because there is no previous discourse, the database pa-meetings (which includes pa-meeting35) are searched. By use of pragmatic rules {6}, which associate the speaker with a meeting as a participant, the possible matches for the pa-meeting in (5) are small and pa-meeting35 can be selected. One experimental hypothesis not fully tested is that for any definite noun phrase in discourse which presupposes an existing object, the database will have a frame, like pa-meeting35, which uniquely matches that noun phrase.

From these examples and the description of reference finding with sub-topics, what remains to be done for PAL to understand examples (6) through (9) can be clearly stated. A notion of parallel topics must be included in the topic recognition and the use of assertions as reasons must be incorporated to discourse scenarios to deal with the complexities in these examples.

The discourses which have been presented are a step along the way to understanding the complexities of reference and speech acts that occur in stories like the one below from Rosenberg [1976].

John murdered his wife. The day before, a load of rat poison was delivered to the house.

It appears that story telling, which is largely assertional and contains sentences that appear unconnected at the surface level, consists of many deep inferences that tie the assertional forms together. To understand

"simple" stories, where one must be able to infer that poison can be a weapon, that John had to have a weapon to commit the murder and so on, the complexities described in this paper must first be resolved. PAL's approach to understanding the example discourses indicates how the first step in the levels of complexity for reference disambiguation and speech act interpretation can be reached.

Acknowledgements I want to thank Gretchen Brown, Ira Goldstein, Beth Levin and Steve Rosenberg for their constructive comments on various revisions of this paper.

Notes {1} The term speech act here harks back to the classic work of Searle[1969] on the definition and use of speech acts.

{2} The referent of *it* cannot be the sentence verb phrase because the verb phrase is want to meet and this cannot be arranged.

{3} I mean here by topic a notion similar to the notion of theme of Kuno [1975]: theme is what the rest of the sentence is about. While Kuno's definition is somewhat vague, his examples are insightful and convinced me that the notion ought to be explicit in my work. I have not retained his term for the notion as topic is more in line with older notions of topic (such as that of Moore [1967]).

{4} Deutsch [1975] first showed that this sort of phenomena occurred in discourse.

{5} The modules include building frames from case frames, resolving intra-sentential reference and pragmatics. See Bullwinkle [1977].

{6} Pragmatic rules used in PAL are rules of commonsense knowledge. They include rules such as "the speaker who requests a meeting to be scheduled is likely to be a participant." Pragmatic rules are discussed in Bullwinkle [1977].

References

Bullwinkle, C.[1977] *The Semantic Component of PAL: the Personal Assistant Language Understanding Program*, MIT AI Laboratory Working Paper, March.

Charniak, E. [1972] *Toward a Model Of Children's Story Comprehension*, M.I.T. A.I. Lab TR-266.

Deutsch, B. [1975] *Establishing Context In Task-Oriented Dialogues*,

Proceedings of the 13 Annual Meeting of ACL, AJCL Microfiche 35.

Goldstein, I. P. and R. B. Roberts. [1977] *NUDGE, A Knowledge-based Scheduling Program*, M.I.T. A.I. Lab memo 405.

Lasnik, Howard. [1976] *Remarks on Co-reference*, Linguistic Analysis, Volume 2, Number 1.

Keyser, S.J. and P. M. Postal. [1976] Beginning English Grammar, Harper and Row, New York.

Kuno, Susumo. [1975] "Three Perspectives in the Functional Approach to Syntax," *Papers from the Parasession on Functionalism*, Grossman et al. (eds.), Chicago Linguistics Society.

Moore, T. [1967] *The topic-comment function: a performance constraint on a competence model*. University of California at Los Angeles dissertation.

Rosenberg, S. [1976] *Discourse Structure*, Working paper 130, M.I.T. lab.

Rosenberg, S. [1977] *Frames-based Text Processing*, paper submitted to IJCAI-77.

Ross, J.R. [1967] *Constraints on Variables in Syntax*, unpublished MIT doctoral dissertation, Cambridge, Massachusetts.

Searle, J.R. [1969] Speech Acts An Essay in the Philosophy of Language, Cambridge University Press.