ICOT Technical Report: TR-376

TR-376

Preference Judgement in Comprehending Conversational Sentences Using Multi-Paradigm World Knowledge

by
T. Ukita, K. Sumita, S. Kinoshita,
H. Sano and S. Amano(Toshiba)

May, 1988

ICOT

Mita Kokusai Bldg. 21F
4-28 Mita 1-Chome
Minato-ku Tokyo 108 Japan

(03) 456-3191~5
Telex ICOT J32964

Institute for New Generation Computer Technology

PREFERENCE JUDGEMENT IN COMPREHENDING CONVERSATIONAL SENTENCES USING MULTI-PARADIGM WORLD KNOWLEDGE*

Teruhiko UKITA, Kazuo SUMITA, Satoshi KINOSHITA, Hiroshi SANO , and Shin'ya AMANO**

Toshiba Corporation Research and Development Center
Komukai-Toshiba-cho 1, Saiwai-ku, Kawasaki 210, Japan

ABSTRACT

In interpreting user's utterances, resolving ambiguities is one of the most fundamental problems to overcome in the development of a question-answering system. This paper describes a framework for understanding conversational language, using the multi-paradigm knowledge representation ("frames" and "rules") of an objective field. Based on the knowledge representation system, the procedure judges preference for candidates of interpretations by identifying causal relationship with messages in the preceding context, where the causal relationship is used to supplement some shortage of information and to give either an affirmative or a negative explanation to the interpretation. The procedure has been implemented in an experimental question-answering system, whose current task is consultation in operating an electronic device (VTR).

1. INTRODUCTION

In building a question-answering system accepting natural language input, it is necessary to deal with a problem wherein a hearer should identify concepts expressed by a user. From the beginning of research on a question-answering system [Winograd 72], this problem of identifying concepts has been studied as anaphoric references and elliptic expressions. Determining anaphora and complementing ellipses represent one of the most fundamental problems which prevent understanding a user's message.

* This work is supported by ICOT(Institute for New Generation Computer Technology).

** Currently he works with ICOT.

To identify concepts regarding objects and events correctly, the machine system should acquire knowledge on an objective task; "world knowledge". Single-paradigm knowledge representation, such as rules, predicates, or frames, is not efficient enough to represent both events and objects. Although a multi-paradigm knowledge representation, for example a combination of frame and rule, might be suitable for representing both linguistic and extra-linguistic knowledge, it has not been fully discussed how world knowledge should be represented and used in language comprehension process. The knowledge representation system in language understanding should represent concepts necessary to accomplish the objective task and, at the same time, to represent some linguistic knowledge required to understand an input sentence. In addition to this so-called "static" knowledge, a knowledge processing system should be capable of managing some information that a user presents in the course of conversation. This capability to manage contextual information is not clearly formulated.

From a linguistic point of view, identifying a user's concept is a basic and essential problem in resolving anaphoric reference and elliptic expression. The anaphora and ellipsis problems have been analyzed and managed, from various points of view (e.g., [Webber 80]), including syntactic and semantic constraints, an inferential approach to find antecedents, and the use of focus to resolve discourse ambiguity. Though focus is a comprehensible concept, to explain anaphoric reference, its extracting process is not well-established. Sidner proposed a "bootstrapping" procedure, which decides focus by verifying its consistency with the contexts [Sidner 83]. In her approach, the inference function is used to verify consistency of the focused object with the contexts. It can be seen that extracting an appropriate focus from the discourse context is a process involved in understanding the discourse itself.

This paper reports means that the authors have explored using a multi-paradigm knowledge representation system, to resolve conceptual identification ambiguity. In the subsequent sections, principles to resolve anaphoric ambiguity are discussed and a practical procedure, using a multi-paradigm knowledge representation system, are described. Then, some trials using an experimental system, whose current task is consultation for using electronic equipment (VTR), will be reported.

2. PRINCIPLES IN COMPREHENDING INPUT SENTENCES

The problem involved in resolving anaphoric ambiguity can be viewed as a selection of the most plausible interpretation corresponding to a new input sentence. In this process, the inference function is required to decide the referents, as well as to extract implicit antecedents. This section focuses on principles to guide a procedure to disambiguate an anaphoric referent with a knowledge base from the viewpoint of designing a practical procedure to interpret a new input sentence.

Research efforts have been carried out in regard to various principles and algorithms used in understanding discourse. After Winograd showed a heuristic way to use knowledge in understanding input sentences [Winograd 72], Rieger arranged inference categories and proposed a reference establishment method in a heuristic way [Rieger 75]. "Coherency" had been proposed as a general concept, to explain discourse structures [Hobbs 77; Grosz,et al. 83], and various phenomena in discourse, including anaphoric identification, are discussed in accordance with this concept. Joshi and Weinstein proposed "appropriateness", which is defined by the number of common objects between sentences, to evaluate magnitude of coherency [Joshi and Weistein 81]. Though this notion explains relations among input sentences, it seems insufficient for designing a practical procedure to resolve ambiguities in concept identification, especially to resolve anaphoric ambiguities for sentences including negative expressions.

A psychologist proposes general principles of understanding sentences as a human mental process [Uchida 82]; they are, consistency of a new interpretation with hearer's knowledge, optimality of the new representation, and open-endedness of the knowledge system. From the viewpoint of designing a machine-oriented procedure, let us consider the former two terms, which seem to be of great importance in resolving ambiguities of concept identification.

Consistency is the most fundamental property that the new interpretation should have. The syntactic and semantic constraints can be viewed as a basic filter, which protects us from adding a superficially contradictory representation. In general, the newly added interpretation should not violate consistency with the speaker's previous utterances, as well as with the knowledge base of the machine system; this may be assumed at least within the scope of a simple question-answering system. At the same time, the notion of consistency claims that the new possible interpretation has some relevance with the previous user's messages and/or the hearer's knowledge base, since there might be a meaningless interpretation that is completely irrelevant, but still 'non-contradictory' to its context. When this assumption is posed, the machine system can exclude tentative interpretations which are contradictory to previous messages and/or the built-in knowledge base.

Optimality can be viewed as a principle wherein the machine system selects the most preferable interpretation. When there are more than two possibilities for an anaphoric referent, the system must decide which candidate is preferable. This preference judgement should be done during an utterance, since no delay in response is basically permitted, especially in a question-answering system. In general, it can be said that the system should search for an interpretation which can supplement a shortage of information or which can produce effective information; in short, the system should find the most informative interpretation. To do so, the system can use causal relations between events which would be prepared in a knowledge base for the system. In resolving anaphoric ambiguity, though the inference function has been used to find inferential antecedents, it should also be used to decide upon an informative interpretation.

3. KNOWLEDGE AND PROCEDURE FOR COMPREHENDING SENTENCES

3.1 Concept Representation for Task Domain

To comprehend language, a hearer must acquire various concepts in regard to events and objects in a task domain. These concepts are generally broken down into events and objects and, therefore, the knowledge for concepts can be classified into the following three categories.

(a) Knowledge of objects and their relationships

An "object" distinguishes an entity or a thing, with or without physical existence, and roughly corresponds to that denoted by nominals. Features and attributes of an object, that characterize it from other objects, should be described as essential information. These features include relationships with other objects, as well as attributes of the object alone.

In this paper, a frame system [Minsky 75] is employed to represent objects and relationships between them in an efficient way. Figure 1 shows a frame example (a part of a full description for a VTR). This example shows that a concept 'VTR' is a sub-set of an electrical_appliance, and its 'status' attribute can take one (car_num=1) value (stand_by, record, ...). Some attributes of an object can be described in its super-class frame as well as in that object description. The standard slots used in describing electronic equipment are super-class, sub-class, has_part, and some status of the equipment (power status etc.). These relationships between objects are traversed during a process of understanding input sentences to extract referents for anaphoric and elliptic expressions. In this paper, a combination of a frame and accompanying rules (described below), is used to represent an object and some related functions of the object, and is called a "schema".

(b)Relationships between an object and an event

An event distinguishes some kind of action or movement of an object and it is expressed as a case-frame, which can be viewed as a "predicate" with a labeled argument. Figure 2 shows an example description for a verb "osu (push)", where slots (agent, object) have a

```
schema(cls,vtr,
  [(superC,      [(value,[electrical_appliance])]),
   (has_part,    [(value,[power_switch,
                  playback_switch,...])]),
   (status,      [(car_num,1 ),
                  (enumeration,[stand_by,playback,
                        record, ...])]),
   (power,       [(car_num,1 ),
                  (enumeration,[on,off ])]),
   (channel      [(car_num, 1),
                  (value_class, integer),
                  (value_condition, (_,V,
                        (V>=1,=<12))) ])
    ......       ],
  [s_rule(power, power_on,VTR,
       ( kr_schema(VTR,[(power,[off])],true),% if
         event(push,1,[(object,[SW])])  ),
       kr_schema(VTR,[(power,[on])],true),  %then

       [('$var_constraint',(
          SW#power_switch :-
               part_of(SW,VTR) ))]),
   s_rule(power,power_lamp_on,VTR,
       kr_schema(VTR,[(power,[on])],true),
       kr_schema(LP,[(pw_lamp_status,[on])],true),
       [('$var_constraint',(
           LP#power_lamp :- part_of(LP,VTR))}])
    ]).
```

Figure 1 Schema example for VTR (a part).

"Kr_schema" is a predicate that extracts or verifies slot values. Words beginning with a capital letter are variables.

```
schema(cls, push,
 [(superC,    [(value, [action])]}]),
  (agent,     [(value_class, human_being)]),
  (object,    [(value_class, switch)]),
  (obligatory_case,
              [(value, [object])]  )    ], [ ]).
```

Figure 2 Event expression by schema.

special facet describing constraints on values for these slots. The requirement for slot fulfillment, in implementing a task, is denoted by the obligatory-case slot (the agent slot in the figure is not obligatory, because most omitted subjects of action verbs in Japanese can be considered to be a user). These constraints are used to reduce search space in interpreting an input sentence.

In addition to a standard frame hierarchy for objects, a schema has rules in its part. A rule in a schema (called s_rule) describes a function of an object. A s_rule consists of three parts; "if_part", "then_part", and variable constraints (optional). Figure 1 shows two s_rule examples concerning the VTR. The first one, for example, describes that, if the power-switch is pushed when the power of the VTR is "off", then the power status will be "on". Since s_rules describe functions of an object schema to which they belong, they can be viewed as description of the relationship between an event and an object. In the above example, pushing the power-switch is an action related to the VTR.

(c) Knowledge of events and their relationships

Basic super- and sub-class relations among events are described by a frame system as objects. In addition to this hierarchical description, relationships between events, such as a causal relation, should be described. The causal relation is extra-linguistic knowledge about a task domain.

In this paper, if-then rules, independent of the schema description, are used to represent relationships between events, specifically causal relations between events. Figure 3 shows an example of rule descriptions whose syntax is almost the same as that for s_rules. For the current task, rules are used to represent some procedures to operate a VTR and are prepared for every one of its functions, such as play-back, record, power-on, and so on. The example of this figure shows an operation of "playback" of a VTR. This kind of rule is used for two purposes in the system. The first one is to reduce the ambiguity and judges preference for candidates in interpreting the input sentence. The second is to guide operation of a piece of equipment. Rules representing connectivities among eventual concepts are referred to by a procedure for interpreting an input sentence to decide upon the final interpretation.

```
rule( operation, play_back,
  seq(event( put_on, 1, [( object,[PW])]),
      event( insert, 1, [( object, [K]), ( goal, [VTR])]) ,
      event( push, 1, [( object, [SW] )]) ),
   event( play_back, 1, [( object, [K]),( tool,[VTR])] ),
   [('$var_constraint',(
     VTR # vtr,
     K # cassette,
     PW # power :- attribute_of ( PW, VTR) ;
     SW # power-switch :- part_of ( SW, VTR) ) )]).
```

Figure 3 Operation procedure example by rule.

Begin
↓
Anaphora detection
↓
Referent candidate extraction
↓
Non_contradiction test
↓
Optimal candidate selection
↓
End

Figure 4 Interpreting input message process flow.

In addition to those static descriptions of concepts, inference mechanisms are provided by the knowledge processing component. They include forward and backward reasoning, based on the rules and s_rules, and the inheritance mechanism for the frame system. The status of a piece of equipment is simulated by forward reasoning. The user's contextual information is reflected as the status of an inner model for the equipment. This inner "model" of the device enables us to inspect the up-to-date status for the equipment and, therefore, to extract candidates for referent that are not presented explicitly in the user's utterances.

3.2 Procedure for Comprehending Sentences

Figure 4 shows a procedure flow for interpreting an input sentence after syntactic analysis. There are four major steps; anaphora detection, referent candidates extraction, non-contradiction test, and optimal candidate selection, which are performed for each predicate contained in a sentence.

(a) Anaphora detection

As the first step in disambiguating anaphora, the procedure detects anaphoric indicators for a sentence. Pronouns and definite articles, which are explicit anaphoric indicators, are detected. At the same time, nouns themselves are extracted for concept identification. In addition to them, omitted obligatory cases of the predicate in a sentence are detected as ellipses. In a situation involving a question-answering system, pragmatic ellipses should be detected, when an interpretation attained for an input sentence is insufficient to form a command for task implementation. In practice, they are detected through obligatory cases. Distinction between obligatory cases and optional cases is a difficult problem for us to describe and a certain criterion is required to distinguish these two categories. Currently, the authors have decided that an event expression, necessary to accomplish the task and represented in rules, is the basic template for events and that case-slots of that template are marked as obligatory.

```
schema( ins, push#1,
    (class,    [(value, [push])]),
    (object,   [(value, [playback_button#1])]),
    (affirmation,[(value,[yes])]) ).
```

(a) schema representation for "I pushed the play-back button".

```
schema( ins, push#1,
    (class,    [(value, [push])]),
    (object,   [(value, [playback_button#1])]),
    (affirmation,[(value,[no])]) ).
```

(b) schema representation for "I didn't push the play-back button".

Figure 5 Schema representation examples for input sentences.

After detecting the anaphoric and elliptical indicators, the procedure produces a temporary instance schema for every indicator. The procedure then inserts the temporary schema into the slot representing the case of the predicate. At this point, for instance, a sentence "I pushed the play-back button" is represented as shown in Figure 5(a).

(b) Referent candidate extraction

For every anaphoric indicator, the procedure searches for candidates corresponding to the indicator. The procedure searches for "instance schemata", which belong to either the same class as the indicator, a sub-class of the indicator's class, or a class connected to the indicators's class by the 'has_part' link. In searching for candidates, the procedure uses constraints on "word class", written at every slot for an event template that restrict the class of objects, to reduce the number of objects to be searched for. For example, for the case of the word "push" shown in Figure 2, a slot value for a "push" object is restricted to switch (a switch includes a button at a conceptual level) by describing "value_class". When an object for a sentence is omitted, the procedure searches for instance schemata, belonging to the switch class and its sub_classes. In general, the number of new tentative representations for the sentence becomes plural, because most verbs have more than two cases and there are usually plural candidates for every case. This search process is carried out by the knowledge representation system invoked by the procedure.

(c) Non-contradiction test

Non-contradiction test of a possible interpretation to its context is accomplished in two subsequent steps. The first step is a test for non-contradiction within a single sentence (predicate), where the corresponding inner representation has undetermined slot values (case values for a predicate). This step is effective to rule out improper combinations of case values for a multi-meaning verb, which corresponds to more than two concepts. For example, the Japanese verb "ugoku" can mean 'work' and 'move' in English and these different concepts take different types of cases. The result of this test is a collection of conceptual case frames which consist of proper combinations of a predicate and its corresponding case objects.

The second step in this process is scanning non-contradictions between the new hypothetical interpretation and the previous user's messages. In general, the procedure might have to verify non-contradiction between the new interpretation and the entire knowledge base. However, in such a system whose task is consultation or diagnosis, it may be usual that a user's input is contradictory with regard to the knowledge base, because faulty symptoms are abnormal and contradictory to the normal state. Thus, at present, the procedure verifies non-contradiction between a new interpretation and the previous user's messages.

As an example of this step, let's consider the following sentences:

"... I pushed the playback button.
(But) I didn't push it."

Deciding the anaphoric referent for "it" in the second sentence, the procedure excludes a possibility " it = playback button", and inspects other candidates. Figure 5 shows internal representations for these interpretations of possible candidates, which are easily detected as contradiction, because these two representations belong to the same class and have the same entries, except for the value in 'affirmation slot'.

(d) Optimal candidate selection

When more than two interpretation possibilities remain, the procedure tries to find causal links which have some relationships to the interpretations and judge preference. By finding an interpretation which has some causal links with the previous sentences, the procedure decides the most informative interpretation. In inspecting causal connectivities for interpretations, the procedure searches for rules in the knowledge base, which can be classified into two groups; one is a rule which has both previous and current event in the conditional part and another is a rule that has the previous events in the conditional part and the current event in the conclusive part.

Assume that X denotes an event in a tentative interpretation being processed, A is an event or a conjunction of events corresponding to the preceding input sentences, and B represents an inferable, goal event or an event which explicitly denotes objective of an action in the sentence. The procedure searches for rules in the order shown below (~X means a negation of X).

(1) if A then ~X,
(2) if A & X then B,
(3) if A & ~X then B,
(4) if A then X.

The first and the last rules show direct causal links between the previous input interpretation(s) (=A) and a hypothetical one (=X). The others are rules that have conditional interpretations.

The order of rules to be searched is decided so that it might correspond to the magnitude of reducing ambiguity; i.e., explicit negation of a plausible event can be considered to give a hearer more information than a positive interpretation for the event.

Consider an example sequence of sentences, whose preceding context contains a cassette and a practical image (on TV):

" I pushed the VTR eject button.
It didn't come out."

To decide the referent of "it" in the second sentence whose candidates are the cassette and the image, the procedure finds a rule such as follows:

if (a cassette is inserted in a VTR),
and (the eject button is pushed),
then (the cassette comes out).

In this case, the rule is classified as belonging to category (1), which has a negated event for the input sentence in the conclusive part. Finding the above rule, the procedure recognizes the subject of the second sentence as the cassette.

Let's check another example:

"Though I pushed the button, the VTR didn't play back the tape."

Processing "the button" in the first part of the sentence, the procedure recognizes it as a button controlling play back, from various kinds of buttons that the VTR has mounted on its control panel. This is achieved by finding the

user: ビデオテープを再生したい
(I want to play back the video-tape)*

system: ビデオの電源を入れて、ビデオにカセットテープを入れて、再生ボタンを押してください
(Please turn on the VTR power, insert the cassette into the VTR, and push the play-back button.)

user: ビデオにカセットテープを入れて、再生ボタンを押した
(I inserted the cassette in the VTR and pushed the play-back button)*

system: はい、わかりました (OK)

user: 動かない
("It's not working" or "It's not moving")**

system: ビデオの電源はオンですか
(Is the VTR power 'ON'?)

user: はい (Yes)

system: 再生ランプは点灯していますか (Is the play-back lamp lit?)

user: 点灯しています (It's lit)*

system: VHF出力端子はテレビに接続されていますか
(Is the VHF output terminal connected to the TV set?)

user: いいえ (No)

system: VHF出力端子をテレビに接続してください
(Please connect the VHF output terminal to the TV set.)

Figure 6 Dialogue example by the experimental system.

* In Japanese, subjects are usually omitted in such a situation.
** Japanese verb '動く (ugoku)' has multiple meanings.

following rule:

if (VTR power is on),
and (a cassette is inserted into the VTR),
and (the playback button is pushed),
then (VTR plays back the cassette).

In this case, the above rule is used as a category (2) rule (X = 'the playback button is pushed', A = nil, and B = 'VTR plays back the cassette').

After the procedure identifies an anaphoric referent, it merges the identified interpretation into the previous one by replacing the link pointer in the event schema. Then, the procedure goes on to the next predicate of the utterance.

4. IMPLEMENTATION AND CURRENT STATUS

The knowledge representation system and the procedure for comprehending input sentences are implemented on a personal sequential inference machine (PSI-II) developed by ICOT. The knowledge base has been preprared from an operation manual for a VTR by hand. Currently, schemata and rules number about 300 (including 180 s_rules) and 30, respectively. Combined with a Japanese analyzer using an efficient parsing mechnism [Matsumoto 87] and a task execution module, the system can answer user's requests in near real time (typically 5 seconds). The sentence analyzer morphologically and syntactically analyzes an input string ("kana-kanji" representation) and outputs a connected string of case frames. Then the discourse analysis procedure described in this paper decides the most plausible interpretation for the input sentences. Finally the task execution module analyzes the structure of the input interpretation and tries to resolve the user's request using the knowledge processing module (therefore, in the experimental system, the knowledge processing module is used for both discourse analysis and problem solving). A current dictionary for the parser consists of about 1000 words (roughly about 700 nouns and 300 verbs), most of them appear in the operational manual. Figure 6 shows a conversation example using the experimental system.

A preliminary test for the procedure described in this paper was performed, in which a few persons talked the system to guide operations of the VTR. For more than 200 occurences of anaphoric and elliptic expressions, about 80 percent expressions were correctly resolved. In the correctly identified expressions, eight instances (5%) were successfully resolved by the preference judgdement by the procedure described in this paper. Most mis-identifications occurred for predicate ellipses and implicit object antecedents. The former was mainly caused by the limitation of the procedure, since the current system can complement an elliptic predicate for a just prior explicit predicate. On the other hand, most of the latter mis-identifications were caused by a lack of knowledge on causal rules for status maintenance for the VTR (which should be prepared as s_rules).

5. CONCLUDING REMARKS

The framework for understanding conversational sentences, which uses world knowledge to disambiguate concept identification, in both positive and negative ways, provides a firm basis which facilitates the extraction of valuable information contained in the described input message. The knowledge system, with multi-paradigm representation, is used to represent concepts of objects and events in an objective task domain. Based on this knowledge base, the procedure for comprehending input sentences resolves ambiguities in correspondence between user's notification and system's concepts.

The procedure and the knowledge representation described in this paper show an inferential approach to comprehending conversational sentences. In constrast to an ordinary use of causal relations to fill gaps between sentences [e.g., Hobbs 79], the procedure in this paper ranks plural candidates by inspecting relationships between events, which are prepared in the form of rules. One of the features of the procedure is that it uses the rules to verify the relationship between events, including a negative expression: It orders the candidates for interpretation by considering information quantity for a hearer (a machine system). This idea might have its basis on the concept that a speaker would not utter a message which is obvious or obviously inferable by the hearer.

On acquisition of concepts in an objective task, it should be noted that knowledge representation for an object and an event must be prepared *a priori*. In the experimental system, an operation manual for the equipment is used as material from which to extract basic objects, events, and some procedures to operate the device. Multi-paradigm representation has an advantage in efficient representation for these concepts, which appear in the manual.

There is a major problem remaining for further study. At present, the procedure cannot accurately identify ellipsis for a predicate and anaphora for a chunk greater than a sentence referent. For these problems, a method to gather context information together is required. The authors are now exploring a way to deal with this.

REFERENCES

[Grosz, et al. 83] Grosz,B.J., Joshi,A.K., and Weistein, S., "Providing a Unified Account of Definite Noun Phrases in Discourse", Proc. 21st Annual Meeting of the Association for Computational Linguistics, pp.44-50 (1983).

[Hobbs 77] Hobbs,J.R., "Coherence and Interpretation in English Texts," Proc. IJCAI-77, pp.110-116 (1977).

[Hobbs 79] Hobbs,J.R., "Coherence and Coreference," Cognitive Science, Vol.3, pp.67-90 (1979).

[Joshi and Weinstein 81] Joshi,A.K. and Weinstein,S., "Control of Inference: Role of Some Aspects of Discourse Structure -Centering", Proc. IJCAI-81, pp.385-387 (1981).

[Matsumoto 87] Matsumoto,Y. and Sugimura R., "A Parsing System Based on Logic Programming," Proc. IJCAI-87, pp.671-674 (1987).

[Minsky 75] Minsky,M., "A Framework for Representing Knowledge", in P.H.Winston (Ed.), The Psychology of Computer Vision, NewYork : McGraw-Hill (1975).

[Rieger 75] Rieger,C.J., "Conceptual Memory and Inference", in R.C.Shank (Ed.), Conceptual Information Processing, North-Holland / American Elsevier (1975).

[Sidner 83] Sidner, C.L., "Focusing in the Comprehension of Definite Anaphora", in M.Brady and R.C.Berwick (Eds.), Computational Models of Discourse, The MIT Press (1983).

[Uchida 82] Uchida, N., "Understanding Sentences and Knowledge", in Y.Saeki (Ed.), Inference and Understanding, Tokyo Daigaku Shuppan (1982) (in Japanese).

[Webber 80] Webber, B.L., "Syntax beyond the Sentence: Anaphora", in R.J. Spiro, et al. (Eds.), Theoretical Issues in Reading Comprehension, Lawrence Erlbaum Associates, Inc. (1980).

[Winograd 72] Winograd, T., "Understanding Natural Language", Academic Press (1972).