

Bolt Beranek and Newman Inc.



AD A092971

LEVEL III

18

Report No. 4513

14 BBN-4513

1084715

SC

6
**Research in Knowledge Representation
for Natural Language Understanding.**

9 Annual Report.

1 September 1979 — 31 August 1980

15 N00014-77-C-0378
✓ ARPA order-3414

11 Nov 80

12 40

DTIC
ELECTE
S DEC 17 1980
E

Prepared for:

Advanced Research Projects Agency

10 William A. Woods
Ronald J. Brachman
Robert J. Bobrow
P. Cohen
B. Goodman

DDC FILE COPY

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

060100

80 12 15 126

JOB

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER BBN Report No. 4513	2. GOVT ACCESSION NO. AD-A092971	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) RESEARCH IN KNOWLEDGE REPRESENTATION FOR NATURAL LANGUAGE UNDERSTANDING - Annual Report 1 September 1979 - 31 August 1980		5. TYPE OF REPORT & PERIOD COVERED Annual Report 9/1/79 - 8/31/80
7. AUTHOR(s) Woods, W., Bates, M., Brachman, R., Bobrow, R., Cohen, P., Goodman, B., Israel, D., Schmolze, J., Sidner, C.		6. PERFORMING ORG. REPORT NUMBER BBN Report No. 4513
9. PERFORMING ORGANIZATION NAME AND ADDRESS Bolt Beranek and Newman Inc. ✓ 10 Moulton Street Cambridge, MA 02238		8. CONTRACT OR GRANT NUMBER(s) N00014-77-C-0378 ✓
11. CONTROLLING OFFICE NAME AND ADDRESS Office of Naval Research Department of the Navy Arlington, VA 22217		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		12. REPORT DATE November, 1980
		13. NUMBER OF PAGES 33
		15. SECURITY CLASS. (of this report) Unclassified
		16a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Distribution of this document is unlimited. It may be released to the Clearinghouse, Department of Commerce, for sale to the general public.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Cognitive science, artificial intelligence, natural language understanding, parsing, syntax, semantics, speaker meaning, knowledge representation, semantic networks, KL-ONE, parallel algorithms, anaphora interpretation, ATN grammars, decision making, vague requests.		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) BBN's ARPA project in Knowledge Representation for Natural Language Under- standing is aimed at developing techniques for computer assistance to a decision maker in understanding a complex system or situation using natural language control of an intelligent graphics display. The work that we have been doing falls into three classes: fluent natural language under- standing in a graphics context - including helpful systems that go beyond mere passive execution of literal instructions, fundamental problems of		

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

Unclassified.

cont'd.

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

Unclassified.

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

20. Abstract (cont'd.)

knowledge representation and use, and abstract parallel algorithms for knowledge base inferential operations. In this report, we will give a brief summary of the activities of this research project during the past year, in particular in the areas of research on parallel algorithms and VLSI, research on the KL-ONE system, and research on natural language understanding. In addition, we document publications written, presentations given and workshops attended.



Unclassified.

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

Report No. 4513

**RESEARCH IN KNOWLEDGE REPRESENTATION
FOR NATURAL LANGUAGE UNDERSTANDING**

Annual Report
1 September 1979 to 31 August 1980

Principal Investigator:
William A. Woods
(617) 497-3361

Accession For	
NTIS GRA&I	☒
DDC TAB	☐
Unannounced	☐
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or special
A	

Prepared for:

Defense Advanced Research Projects Agency
1400 Wilson Boulevard
Arlington, VA 22209

ARPA Order No. 3414

Contract No. N00014-77-C-0378

Effective Date of Contract:
1 September 1977

Contract Expiration Date:
30 September 1981

This research was supported by the Advanced Research Projects Agency of the Department of Defense and was monitored by ONR under Contract No. N00014-77-C-0378. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

TABLE OF CONTENTS

	Page
1. INTRODUCTION	1
2. RESEARCH ON PARALLEL ALGORITHMS AND VLSI	5
3. RESEARCH ON THE KL-ONE SYSTEM	7
4. RESEARCH ON NATURAL LANGUAGE UNDERSTANDING	13
4.1 Parsing and Semantics	13
4.2 Pragmatics of Discourse	18
4.3 Interpreting Vague Requests	22
5. CONCLUSION	25

1. INTRODUCTION

BBN's ARPA project in Knowledge Representation for Natural Language Understanding is aimed at developing techniques for computer assistance to a decision maker in understanding a complex system or situation using natural language control of an intelligent graphics display. The motivating need is that of a military commander in a command and control context both in strategic situation assessment and in more tactical situations - especially in crisis situations. The underlying assumption of this work is that in a crisis situation the commander needs an extremely flexible system, capable of manipulating large amounts of data and presenting it on a graphical display in a variety of ways until the commander feels satisfied that he has a grasp of the situation. Such a system would have abilities to display many kinds of different map overlays, an ability to change the kinds and amounts of detail shown, an ability to conveniently construct unique kinds of displays to suit the situation at hand, as well as the ability to display tabular and graphical information and present textual material in ways that are easily comprehensible.

Techniques to produce such displays on demand, in response to high level specifications of what they should contain, do not currently exist and will require significant breakthroughs in

areas of language understanding, knowledge representation, and knowledge based inference. The work that we have been doing falls into three classes, successively motivated by the initial goal of providing powerful computer assistance to a commander in a complex decision-making task. These areas are: fluent natural language understanding in a graphics context - including helpful systems that go beyond mere passive execution of literal instructions, fundamental problems of knowledge representation and use, and abstract parallel algorithms for knowledge base inferential operations.

The major accomplishment in this work so far has been the development of the knowledge representation system KL-ONE and its use in the construction of an experimental prototype system that understands English requests for display manipulation. This system parses and interprets English requests, synchronized with pointing events on a screen, and produces appropriate display actions on a bit map graphics display in response. It permits a user to request portions of a display to be shown, objects in the display to be made visible or invisible, attributes of objects pointed to be displayed, and specification of refocusing requests by means of statements of constraints on what is to be visible.

The knowledge representation system KL-ONE is used in this system to organize the semantic interpretation rules used to interpret sentences, to organize the models of the user's goals and beliefs (which are used to fill in details that are not explicit in the input), and to organize the knowledge of displays and display forms that are used to draw the pictures on the screen. The knowledge structuring capabilities of the KL-ONE system have proven themselves very powerful in this system, and the extent to which the same structures have proven useful in qualitatively different parts of the system gives evidence of the robustness of these capabilities.

In addition, a component of this project is devoted to cooperation with the ARPA sponsored Consul group at ISI, to provide them with current versions of the RUS parsing system and the KL-ONE knowledge representation system and to work on knowledge representation problems that arise out of that work. This interaction has so far been fruitful for both groups. A similar cooperation exists between the KL-ONE group and the ARPA supported AIPS project [Greenfeld and Yonke, 1979] [Zdybel, Greenfeld and Yonke, 1980] at BBN, which is also using KL-ONE.

KL-ONE currently has an exceptionally good representation for the inheritance relations among structured concepts, including the relations between corresponding parts of their

structures. However, there are many subtleties of representation that are still undergoing active investigation as part of the knowledge representation effort and that require continued development. For some recent references, see Brachman [1979], Brachman and Smith [1980] and Brachman et al. [1979].

In this report, we will give a brief summary of the activities of this research project during the past year. More detailed accounts of activities will be contained in technical reports that will be published from time to time. This report will present summaries of activities in the areas of research on parallel algorithms and VLSI. research on the KL-ONE system, and research on natural language understanding.

2. RESEARCH ON PARALLEL ALGORITHMS AND VLSI

A significant component of our research in knowledge representation concerns the development of algorithms that will operate on KL-ONE representations to produce the kinds of inference that are important to knowledge based applications. One of these algorithms is an operation we have called MSS (for most specific subsumer). The algorithm takes a description and locates the most specific concepts in a KL-ONE network that subsume it and should be connected to it as superconcepts. This algorithm can be used as the basis for a taxonomically oriented, rule-based system whose operation is to classify a given situation in a conceptual network of situations and thereby find a concisely indexed set of rules to be applied to that situation.

During the past year, we have made some improvements to our MSS algorithm and have also begun the exploration of parallel algorithms for performing this operation on abstract parallel machine architectures. Previously, we developed a formal language for specifying parallel marker passing algorithms [Woods, 1979] and have specified a version of the MSS algorithm in this language. Now we have implemented an initial version of a simulator for this abstract architecture and begun the process of testing it with the initial MSS algorithm.

Also during the past year, we have been exploring the possibilities of VLSI architectures for realizing the kinds of parallel algorithms that we are developing. Pursuant to this end, two members of the project completed a special course at BBN on VLSI design. A VLSI chip was designed that implements a critical component of a proposed design for a special architecture ATN grammar machine. This chip, called a contention gate, provides fast and efficient automatic contention resolution when multiple hypotheses in a search operation want to communicate to a central hypothesis dispatcher. This chip was included in the MPC580 multiproject chip.

In the area of abstract algorithms, we have also been developing several abstract automata that are generalizations of ATN's, one of which is a way of viewing KL-ONE networks as a kind of hierarchy of ATN's with inheritance [as was presented in Woods, 1978]. A paper describing some of this work, [Woods, 1980] appeared this year in the *American Journal of Computational Linguistics*.

3. RESEARCH ON THE KL-ONE SYSTEM

This year saw a surge of research on the abstract principles of KL-ONE, as well as being a period of renewed effort on the implementation and maintenance of the KL-ONE system.

The major focus of research this year in the abstract conception has been in clarifying the semantic impact of some of the basic object types. In particular, a detailed study of the SuperC Cable revealed that it had at least two possible intents - as a simple universal statement (as in "all camels are mammals"), and as a definition forming operator (e.g., defining the Concept, TRIANGLE, in terms of a more general Concept, POLYGON). The former type of statement is actually just one of many types of assertions we would like to make, and thus we have chosen to develop a general mechanism for representing statements, and allocate to the SuperC Cable purely definitional meaning.

The insistence on definitional commitment in the SuperC link at first appears somewhat severe, since much of our "knowledge about the world" is drawn from observation of entities naturally occurring in the world. However - provided that we augment it with a facility for representing "natural kinds" and making contingent statements - a definitional structure-building facility gives a clean semantics, and the ability to define composite Concepts, both of which tend to be lacking in most

knowledge representation languages.

Work this year has proceeded from this conclusion, and we have slowly gotten a feel for what it means for Concept structures to be definitional. Much of what Concepts could previously be used for is now the responsibility of KL-ONE's "Nexus" mechanism. All statements to be made that are not necessarily true (i.e., true in all possible worlds) are to be made with respect to a Context - the representation of a possible world. For example, attributions of descriptions to particular individuals are not to be done with multiple SuperC's on Individual Concepts, but rather by attachment of multiple Concepts (either Generic or Individual) to Nexuses.

It is the Nexus and Context part of KL-ONE that is now under the most intense scrutiny. In research in the last several months, we have developed a proposal for representing arbitrary propositions as intensional objects, and providing a way to assert them in particular Contexts. We also see the need to be able to represent someone's believing a proposition (or some person A believing that some other person B believes a proposition) independent of that proposition's actual truth-value. We plan to provide for representing these beliefs using Contexts as well.

As reported in our semiannual report for 1980, during the

first part of this year, three implementation efforts were undertaken, namely a substantial speedup to the KL-ONE system, a re-organization of the system files and functions, and a set of experimental functions for exploring the use of the different kinds of "meta" relationships in KL-ONE. This work being completed, we concentrated our software effort in the second half of the year on four themes:

1. building packages to assist our implementation work;
2. extending the functional interface to KL-ONE;
3. documenting this interface;
4. fixing software bugs.

The work on utility packages included building new programming aides as well as expanding existing ones, among these an INTERLISP cleanup package designed for the particular environment of KL-ONE development work. In addition, we have expanded somewhat our facility for generating documentation directly from software.

Recent extensions to the KL-ONE interface were aimed at providing more functionality for users. As before, we discriminate between information which is actually stored at a place in the knowledge base and information which is inherited at a place. In several cases, obvious functions for local access that had previously been omitted from the interface were added.

Additionally, some extensions involved filling gaps in the general functionality of KL-ONE. In particular, functions manipulating "values", "value restrictions", and "value descriptions" were completed.

This past year, Role differentiation and modification were improved as well. A detailed examination of the intent of RoleSet differentiation has led to a clearer interpretation of its semantics. In particular, we now see differentiation as somewhat analogous to the SuperC relation, and view it as giving significant internal structure to a RoleSet. Previously, subRoles of a RoleSet were treated more as parts of a Concept than as parts of the Role that they differentiated. We have begun a series of improvements in the set of KL-ONE functions for manipulating RoleSets, among them distinguishing specializations of a Role (its modifications only), from manipulating its subRole descendants (differentiation).

In the latter part of the year, we also began to examine the entire set of functions in the KL-ONE interface to insure that the documentation for each was correct. This process is not currently complete and will continue into the new year. Upon its completion, we will be able to generate complete up-to-date documentation of the KL-ONE system function directly from the software.

Report No. 4513

Bolt Beranek and Newman Inc.

Finally, during the past year, as the number of KL-ONE users grew, so did the number of reports of software bugs. We have tracked down a number of these problems and have corrected them.

Bolt Beranek and Newman Inc.

Report No. 4513

4. RESEARCH ON NATURAL LANGUAGE UNDERSTANDING

Research on natural language has been pursued in three basic areas: parsing and semantics; pragmatics of discourse; and interpreting vague requests. We will discuss each below.

4.1 Parsing and Semantics

In parsing, the work on the RUS Parser has consisted of two broad classes of activities - development of a new, more flexible and space-efficient control structure for the grammar compiler, and redesign and tool-building work designed to make the parser readily maintainable and exportable to external groups such as ISI.

In the original version of the Burton grammar compiler, an entire ATN is converted into a single (generally extremely large) INTERLISP function. The size of the function created for the RUS Parser grammar was so large that it was virtually impossible to compile it, and the interpreted version was too large to fit in a system that had KL-ONE loaded. During this year the grammar compiler was redesigned to allow separate sections of an ATN to be translated into independent functions, and provide an interface that allowed these functions to be run as coroutines so that variations in the ATN control flow could be achieved with suspended alternatives and non-deterministic backtracking. At

the same time, various scheduling primitives that had been added to the system to provide increased determinism in parsing were abstracted and generalized to provide a very flexible scheduling structure. This structure makes it possible to implement such alternatives as selective modifier placement by simply changing one control function and a small number of global variables.

The result of these changes is that the new version of the RUS parser is more efficient than previous versions, its grammar is more readable and understandable, and there is greater flexibility for experimenting with alternative control strategies. It is also possible to modify the grammar quickly even though it is a fully compiled system (i.e., one in which the grammar is in the form of a set of block-compiled INTERLISP functions). It is also possible to load the new system into a version of KL-ONE in INTERLISP-10, without running out of space or using special loading procedures and minimal subsets of the INTERLISP-10 system, as was previously necessary.

In our maintenance work on RUS, the overall structure of the set of files used to work with the RUS parser and grammar compiler was completely revised, so that it is possible to load just that subset of the system which is needed for a given application such as production or grammar development. Standardized loading procedures were developed and automated,

with appropriate on-line prompting and documentation. In order to guarantee that revisions to the grammar and control structure did not reduce previous performance, a standardized test procedure was developed including a database of sentences to be parsed. The grammar was compared with the SRI DIAGRAM [Robinson, J., 1980] grammar, the LUNAR grammar [Woods et al., 1972], Marcus' PARSIFAL [Marcus, 1977] and several other systems, and the differences in syntactic coverage were determined. The earlier version of the RUS grammar covered virtually all the constructions in these systems, with the primary exception of the comparatives permitted in DIAGRAM. The grammar was modified to cover the range of comparatives, and was also extended to include a wider range of verb phrase complementation structures than the other systems. Work is underway to check that the test database in fact exercises all the features of the grammar. This will make it possible to use the test database as part of the documentation of the properties of the grammar and the use of each of its arcs.

In order to make the RUS parsing system more available to external users, it was necessary to make extensive changes and additions to the functions that maintain and access the dictionary. Originally, when the parser encountered a word it did not know, it asked the user to enter its dictionary definition and the user had to be very knowledgeable about the

information contained in that definition and its format. A more helpful function has been written which asks the user a series of questions and provides examples and help if necessary to elicit the information. As a result, the user no longer has to be aware of the internal dictionary structure. Several other functions have been written to facilitate the production of a new dictionary and to allow several dictionary files to be used simultaneously by the parser.

A feature of the old version of RUS provided for fetching words from the dictionary file without all the dictionary definitions being loaded in at once. As INTERLISP changed, it was not possible to use the features of INTERLISP that provided this dictionary file capability, so we developed new functions to do this work. We chose to use the LispUsers hash file package together with a set of utility functions called PropFile specifically for saving property list information on hash files. As another efficiency measure, we are converting the internal representation of dictionary entries from property list format to hash arrays because information can be retrieved from hash arrays about twice as fast as from moderately long property lists.

The work on semantics and knowledge representation consisted of two major parts. First, the properties of the recognition algorithm used in the PSI-KLONE syntactic/semantic interface were

analyzed to understand the demands that it placed on the underlying representation (i.e., the features that would have to be present in KL-ONE in order to represent all the information needed for recognition in the declarative formalism). Second, a new semantic representation has been initially designed, which replaces the LUNAR style extended quantificational system. The new representation distinguishes between the components of meaning in a sentence that delineate the descriptive portion of the sentence (i.e., what type of event/relation/action the sentence expresses) from the combinatoric portion of the sentence (i.e., that information that is normally carried by determiners and was captured in part by the extended quantifiers used in LUNAR). We expect the separation of these two aspects of meaning to facilitate anaphor resolution and other speech act processing.

This work has led to an understanding of an important use of the meta-description facility in KL-ONE, and constraints on its future development. These two areas of research are presented in a brief form in papers published as part of the proceedings of the AAAI (American Association for Artificial Intelligence) conference, and the CSCSI (Canadian AI Society) conference (see Bobrow and Webber [1980a, 1980b]). A more detailed paper is currently being prepared for submission to the journal *Artificial Intelligence*.

4.2 Pragmatics of Discourse

The phrase "pragmatics of discourse" is a general term which describes our work on natural language phenomena concerning the interpretation of reference and the understanding of speaker meaning. Reference phenomena include first reference using definite noun phrases or names, anaphoric reference and deictic reference (i.e., reference by pointing). Speaker meaning includes the use of sentences in a discourse to accomplish some purpose; in the past we have spoken of such uses as "speech acts". Our change of terminology reflects a change in approach which will be discussed further below.

During the past year we have concentrated effort on reference on two problems. First, based on the work of Sidner [1979], a machine for interpreting focus was designed as a finite state machine and was coded for use in interpreting anaphoric expressions. This machine provides the basis for the focusing approach to anaphora resolution; the approach is based on the definition of this machine and an interpreter which uses the machine's choice of focus in a series of rules for choosing an anaphor's interpretation. Encoding of the rules in INTERLISP for the anaphor interpreter remains to be done. Second, in a study of several texts, we initially explored how authors use reference in texts where pictures are presented and included in the

explanation of some object. Since little research has been directed at how people organize their discussions of objects in pictures, this study provides us with some direction for our work in understanding reference in a discussion involving a graphics display.

Also during this year, we initiated a protocol collection, which has provided us with considerable data about how people talk about objects on a graphics display which they are manipulating. Our research plan called for gathering of protocols of user behavior in simulated discourses involving visual display. We designed a set of instructions for protocol collection on two tasks, both involving visual display and natural language communication between a subject and a simulated NLU system (a person was used for the simulation). The two tasks were the design and layout of logic circuits, and the construction of KL-ONE networks for a database of new information. The protocols were collected on three subjects for each task. A report explaining the devices used to collect the protocols, as well as the full body of the protocols, both visual and typescript, is available as an internal report.

The protocols reflect a wide range of behavior, but they are particularly helpful for us for the range of behavior involving speaker meaning, referential phrases, anaphoric phrases and

deixis. Since the protocols were designed to allow speakers to use "free flowing" English, and since a graphics display system complete with a simulated "mouse" was available to the subjects, speakers directed the "system/person" with a variety of referring expressions and with a variety of sentence command, request and comment types. We are using these protocols extensively for our work on speaker meaning and will also use them for our theory of speaker reference.

Finally, our work on speaker meaning has proceeded in two phrases. Earlier this year, we tested the current prototype to eliminate remaining bugs and inconsistencies as well as to extend the demo to include a means of interpreting "where is x?" type questions. The result of these efforts was not only a cleaner prototype but also an understanding of ways to improve upon the interpretation of speaker meaning.

We have begun redesigning our model of speaker meaning along two dimensions. First, we observed the need to allow a greater degree of interaction between hypotheses about the propositional content of the utterance (and the context in which it occurs) and hypotheses about the nature of the speech-act being performed. The central idea in this regard is that the audience is attempting to come up with an explanation of the speaker's utterance (assuming that it was intentional and directed at the

audience), an explanation in terms of beliefs and desires it is reasonable to attribute to the speaker. In short our feeling is that the current design is too bottom-up, too much driven by the utterance-type itself. Part of our redesign includes an expansion of the role of the speaker's plan in the development of the hearer's explanation. The new model infers much more about the speaker's plan, and the speaker's knowledge of the hearer's capacities. In our model, such knowledge is used by the hearer to explain exchanges in which misinformation and miscommunication occurs between conversational participants.

Two papers are in progress on this work. The first, "A proposal for Interpreting Speaker Meaning," is being co-authored by C.L. Sidner and D.J. Israel. A second paper, by Sidner, "Why People Talk That Way," explores the role of speaker meaning in discourse. This paper shows that speakers make use of implicit information, commonly called "speech acts," to make conversation more efficient and to relegate responsibility for certain decisions to their intelligent conversational partner.

A second part of our design requires a clearer analysis of propositional attitudes. In particular we have studied the problems of opacity and of *de re* and *de dicto* belief¹. It is

¹ These terms name classical philosophical distinctions that turn out to be significant for the kinds of knowledge representations we are attempting.

quite clear that these problems must be dealt with in an analysis of planning situations and in modeling one conversational participant's beliefs about the other participant, since the task of our system is to propose and reason about hypotheses involving the imputation of plans and planning behavior to the user. This work has also raised questions about the expressive adequacy of KL-ONE and hence has strongly interacted with work on the abstract principles of KL-ONE.

4.3 Interpreting Vague Requests

One of the central issues in a natural language question answering environment is how to convert 'vague' and 'complex' requests into a sufficiently well-defined internal representation that the computer can work with. A well-defined request is one where the problem statement itself coupled with the system's world knowledge facilitates immediate retrieval of some solution. Vague requests are incompletely specified, i.e., certain of the information needed to answer the question is missing. Complex requests are ones that go beyond a simple request for a single direct retrieval of information.

To consider these issues, we are investigating the interpretation of requests about report generation of business-related activities such as payroll, budget and cost

accounting. The basic constructs used in report generation are charts and tables; the basic constructs of business-related activities are employees, departments, salaries and costs; these objects and actions have been described in KL-ONE. While the actual procedures needed to automatically construct these objects, and perform these actions have not been described in KL-ONE, a delineation of some of the procedures involved has been performed, in particular for: display-table, fill-table, add-row, delete-row, add-column, delete-column, move-row, move-column, system-construct-displayform, system-construct-table, system-generate-report and system-generate-overhead-report.

A syntactic taxonomy of a scenario defining and building a (chargeability) report has also been constructed. For each kind of clause in a sample report generation scenario, a set of semantic rules have been written that map the roles of the syntactic elements into the KL-ONE hierarchy describing the objects and actions in the domain. The current set of rules principally maps roles filled by syntactic elements such as subject, object, and prepositional phrase into the case roles of the elements, such as AGENT, OBJECT, INSTRUMENT, LOCATION, BENEFICIARY, and TIME. Procedures for carrying out the requested actions can be attached to the domain actions.

In the cases where not enough information is given to carry out a particular action or to decide what action is being requested, a "vagueness" resolver has to be employed to try to reformulate the request into a well-specified problem or action. We are currently exploring a number of design issues of such a vagueness resolver. At present, vague requests will be classified in one of two ways: (1) the description of the request is not complete enough to allow the semantic interpretation rules to map out a specific interpretation of the request; and (2) a particular interpretation can be found but is filled with vague values.

A central element of the vagueness resolver will be a matcher that matches (partial) descriptions of a request to elements of the KL-ONE data base. We are currently investigating the kinds of matching capabilities that will be useful for systems using KL-ONE so that the matcher employed by the vagueness resolver can be used for other tasks. Some of the issues under consideration include: (1) possible use of a full-scale theorem prover; (2) specification of circumstances for forced match; and (3) what constraints can be formalized to limit the search for a match. Currently, we believe a heuristic-driven unification matcher with scoring should suffice for the needs of a vagueness resolver.

5. PUBLICATIONS

To conclude this report we present here a list of publications, including abstracts, by members of the research group. Also included are presentations given and workshops attended.

PUBLICATIONS

Brachman, R.J. and Smith, B.C. SIGART Newsletter, Special Issue on Knowledge Representation, No. 70, February 1980.

Abstract

In the fall of 1978 we decided to produce a special issue of the SIGART Newsletter devoted to a survey of current knowledge representation research. We felt there were two useful functions such as issue could serve: elicit a clear picture of how people working in this subdiscipline understanding knowledge representation research and provide a document that would enable to the reader to sense how different research endeavors around the world fit into the field as a whole. This issue presents the result of the survey. We summarize the form the survey took, strategies we used to analyze responses as well as the responses themselves.

Bobrow, R.J. and Webber, B.L. PSI-KLONE: Parsing and Semantic Interpretation in the BBN Natural Language Understanding System. Proceedings of the CSCSI/SCEIO Conference, Victoria, B.C., May 1980.

Abstract

This paper describes the syntactic and semantic processing components of a natural language understanding system currently under development at BBN. There are several interesting features of this system which this paper will highlight. The first is a framework for natural language parsing (called the **RUS** parser) which combines the efficiency of a semantic grammar with the flexibility and extensibility of modular syntactic-semantic processing. The second (the **PSI-KLONE** interface) comprises two descriptive taxonomies represented in the **KL-ONE** formalism [Brachman, 1979] which represent, first the system's knowledge of interpretable syntactic-semantic patterns, and, second, the system's semantic knowledge of possible objects, events and relationships. These taxonomies facilitate the two major tasks of the system's semantic processor:

1. providing feedback to the syntactic processor, and
2. providing semantic interpretations for individual phrases.

A third interesting feature of the system is touched upon only briefly - its treatment of natural language quantification in terms of a combinatoric problem to be solved, to whatever extent necessary, by a pragmatics/discourse component.

Bobrow, R.J. and Webber, B.L. Knowledge Representation for Syntactic/Semantic Processing, Proceedings of the First Annual National Conference on Artificial Intelligence, Aug. 1980.

Abstract

This paper describes the RUS framework for natural language processing, in which a parser incorporating a substantial ATN grammar for English interacts with a semantic interpreter to simultaneously parse and interpret input. The structure of that interaction is discussed, including the roles played by syntactic and semantic knowledge. Several implementations of the RUS framework are currently in use, sharing the same grammar, but differing in the form of their semantic component. One of these, the PSI-KLONE system, is based on a general object-centered knowledge representation system, called KL-ONE. The operation of PSI-KLONE is described, including its use of KL-ONE to support a general inference process called "incremental description refinement." The last section of the paper discusses several important criteria for knowledge representation systems to be used in syntactic and semantic processing.

- - - - -

Israel, D.J. What's Wrong with Non-Monotonic Logic. Proceedings of the First Annual National Conference on Artificial Intelligence, August 1980.

Abstract

In accordance with the standard notion of semantic interpretation, a sentence A is a semantic consequence of a set of sentences G iff every interpretation under which all of the sentences of G are true is one under which A is true as well. It follows immediately that the relation of semantic (logical) consequence is monotonic: if A is a

semantic consequence of G, then it is also a semantic consequence of every superset of G. It has been argued (by, e.g., Reiter, Doyle and McDermott) that, in virtue of the monotonicity of its consequence relation, no standard logic is adequate for modeling the reasoning of an intelligent subject whose knowledge and evidence are limited or incomplete, especially if that subject must make judgments based on what it knows to be incomplete evidence. I argue that their complaint rests on a fundamental misconception of the role of logic in reasoning and that their insights can be accommodated by recognizing that reasoning must indeed be non-monotonic; but that logic had best not be.

- - - - -

Sidner, C.L. Focusing and Discourse. Submitted to Discourse Processes (journal), March 1980.

Abstract

In a discourse, speakers talk about something that is of interest to them. They center their attention on a particular element of the discourse, and they talk about it for several sentences of their discourses. This element is called the focus, and the process by which speakers center is focusing. Focusing is a cognitive process which is active during the interpretation of discourse rather than during the interpretation of isolated sentences. This paper describes a process model of focussing which specifies what syntactic and semantic constraints of language and what real world knowledge are needed to track the speaker's focus through a discourse. The paper illustrates that focusing is a well constrained behavior for speakers, and argues that focusing is a necessary condition for maintaining Grice's maxim of conversation.

- - - - -

Sidner, C.L. Focus for the Interpretation of Pronouns., To be submitted to American Journal of Computational Linguistics, October 1980.

Abstract

Recent studies both in artificial intelligence and linguistics have shown that the comprehension of anaphoric expressions in natural language presents numerous difficulties. In this paper a new approach, based on a theory of the process of focusing on parts of the discourse, will be used to explain the interpretation of anaphora. The concept of a speaker's foci will be defined and it will be shown that they can be used to choose the interpretations of personal pronouns. The rules for choosing interpretations will be stated within a framework that shows:

- o how to control search in inferencing by a new method called constraint checking,
- o how to take advantage of syntactic, semantic and discourse constraints on interpretation,
- o how to generalize the treatment of personal pronouns, to serve as a framework for the theory of interpretation for all anaphora.

- - - - -

Sidner, C.L. Focusing and the Comprehension of Definite Anaphora. Chapter for book, Michael Brady (ed.), MIT Press, to appear, Spring 1981.

Abstract

This chapter investigates the process of focusing as a description and explanation of the comprehension of certain anaphoric expressions in English discourse. An abstract machine for

determining the foci of discourse as well as how they move, is presented. The foci chosen are then used by an anaphor interpreter, which interprets a set of rules based on the foci for finding the proper interpretations of anaphora. The chapter includes extensive discussion of rules of interpretation for personal pronouns and for the demonstrative noun phrases with articles this and that.

- - - - -

Woods, W.A. Cascaded ATN Grammars. *American Journal of Computational Linguistics*, Vol. 6, No. 1, January-March 1980.

Abstract

A generalization of the notion of ATN grammar, called a cascaded ATN (CATN), is presented. CATN's permit a decomposition of complex language understanding behavior into a sequence of cooperating ATN's with separate domains of responsibility, where each stage (called an ATN transducer) takes its input from the output of the previous stage. The paper includes an extensive discussion of the principle of "factoring" - "conceptual factoring" reduces the number of places that a given fact needs to be represented in a grammar, and "hypothesis factoring" reduces the number of distinct hypotheses that have to be considered during parsing.

PRESENTATIONS

Ronald J. Brachman:

"How KL-ONE Facilitates Natural Language Understanding",
University of Massachusetts, 12/10/79.

"A Glimpse into the Representational Features of KL-ONE", Carnegie-Mellon University, 3/11/80.

"Taxonomy, Descriptions, and Individuals: Some Representational Features of KL-ONE", Brown University, 4/22/80.

"Natural Language Reflections of Knowledge Representation Distinctions". Sloan Workshop on Knowledge Representation and Language Comprehension, Brown University, 7/12-7/14/80.

"Recent Advances in Representation Languages", AAAI - First National Conference on Artificial Intelligence - 8/19/80.

Robert J. Bobrow:

"PSI-KLONE", University of Toronto, April 1980.

"PSI-KLONE", CSCSI Conference, Vancouver, April 1980.

"The RUS Parser system", Sloan Workshop on Language Processing and Language Learning, Brown University, June 1980.

David J. Israel:

"Non-monotonic logic and non-monotonic reasoning", Workshop on Formal Methods in Artificial Intelligence (sponsored by the Rand Corporation), Monterey Dunes, CA, August 1980.

Panel on Artificial Intelligence and Philosophy (with Bob Moore, Pat Hayes, and John Haugeland) at the First Annual Conference of the American Association for Artificial Intelligence, Stanford University, August 1980.

"What's Wrong With Non-Monotonic Logic", a talk given at AAAI Conference, Stanford University, August 1980.

William A. Woods:

"Knowledge Representation for Command Decision Making", ONR Lectures on Artificial Intelligence, Washington, D.C., July 23, 1980.

WORKSHOPS ATTENDED

Brachman, R.J.

Formal Methods in Artificial Intelligence, Monterey Dunes Colony, CA, 8/15/-8/17/80.

Expert Systems, San Diego, CA, 8/25 - 8/29/80.

Israel, D.J. and Woods, W.A.

Sloan Workshop on Artificial Intelligence and Philosophy, Institute for Advanced Studies in the Behavioral Sciences, Palo Alto, CA, March 1980.

OTHER RELEVANT PAPERS

Brachman, R.J. "On the Epistemological Status of Semantic Networks", in Nicholas V. Findler (ed.), **Associative Networks - The Representation and Use of Knowledge in Computers**. New York: Academic Press, 1979.

Brachman, R.J., Bobrow, R.J., Cohen, P.R., Klovstad, J.W., Webber, B.L., and Woods, W.A. **Research in Natural Language Understanding Annual Report: 1 September 1978 - 31 August 1979**, Report 4274, Bolt Beranek and Newman Inc., August 1979.

Greenfield, N.R. and Yonke, M.D. **AIPS: An Information Presentation System for Decision Makers**. Report No. 4228, Bolt Beranek and Newman Inc., December 1979.

Marcus, M. **A Theory of Syntactic Recognition for Natural Language**. Ph.D. thesis, MIT, 1977. (Also published by MIT Press, 1979)

Robinson, J.J. **DIAGRAM**. Technical Note No. 205, SRI International, Menlo Park, CA, 1980.

Sidner, C.L. **Towards a Computational Theory of Definite Anaphora Comprehension In English Discourse**, MIT AI Laboratory Technical Report No. 537, August, 1979.

Woods, W.A., Kaplan, R.M. and Nash-Webber, B.L. The Lunar Sciences Natural Language Information System: Final Report. Report No. 2378, Bolt Beranek and Newman Inc., June 1972.

Woods, W.A. Generalizations of ATN Grammars, in Research in Natural Language Understanding, Quarterly Progress Report No. 4, Report 3963, Bolt Beranek and Newman Inc., August 1978.

Woods, W.A. Parallel Algorithms for Real Time Knowledge Based Systems, in Research in Natural Language Understanding, Quarterly Progress Report No. 6, Report 4181, Bolt Beranek and Newman Inc., February 1979.

Zdybel, F., Greenfeld, N., and Yonke, M. Application of Symbolic Processing to Command and Control: An Advanced Information Presentation System, Annual Technical Report. Report No. 4371, Bolt Beranek and Newman Inc., April 1980.

Official Distribution List
Contract N00014-77-C-0378

	<u>Copies</u>
Defense Documentation Center Cameron Station Alexandria, VA 22314	12
Office of Naval Research Information Systems Program Code 437 Arlington, VA 22217	2
Office of Naval Research Code 200 Arlington, VA 22217	1
Office of Naval Research Code 455 Arlington, VA 22217	1
Office of Naval Research Code 458 Arlington, VA 22217	1
Office of Naval Research Branch Office, Boston 495 Summer Street Boston, MA 02210	1
Office of Naval Research Branch Office, Chicago 536 South Clark Street Chicago, IL 60605	1
Office of Naval Research Branch Office, Pasadena 1030 East Green Street Pasadena, CA 91106	1
Office of Naval Research New York Area Office 715 Broadway - 5th Floor New York, NY 10003	1

Naval Research Laboratory Technical Information Division Code 2627 Washington, D.C. 20380	6
Naval Ocean Systems Center Advanced Software Technology Division Code 5200 San Diego, CA 92152	1
Dr. A.L. Slafkosky Scientific Advisor Commandant of the Marine Corps (Code RD-1) Washington, D.C. 20380	1
Mr. E.H. Gleissner Naval Ship Research & Development Center. Computation & Mathematics Dept. Bethesda, MD 20084	1
Captain Grace M. Hopper NAICOM/MIS Planning Branch (OP-916D) Office of Chief of Naval Operations Washington, D.C. 20350	1
Mr. Kin B. Thompson NAVDAC 33 Washington Navy Yard Washington, D.C. 20374	1
Advanced Research Projects Agency Information Processing Techniques 1400 Wilson Boulevard Arlington, VA 22209	1
Captain Richard L. Martin, USN 3311 W. Avenue Newport News, VA 23607	1
Director National Security Agency Attn: R54, Mr. Page Fort G. G. Meade, MD 20755	1