SRI International

NATURAL LANGUAGE ACCESS TO MEDICAL TEXT

Technical Note 240

March 1981

By: Donald E. Walker and Jerry R. Hobbs
Artificial Intelligence Center
Computer Science and Technology Division

Project 1944

An abridged version of this paper will appear in the _Proceedings of the Fifth Annual Symposium on Computer Applications in Medical Care_, Washington, D.C., November 1981. (New York: IEEE, 1981.)

This paper was supported in part by NIH Grant LM 03611 from the National Library of Medicine.

333 Ravenswood Ave. • Menlo Park, CA 94025
(415) 326-6200 • TWX: 910-373-2046 • Telex: 334-486

ABSTRACT

This paper describes research on the development of a methodology for representing the information in texts and of procedures for relating the linguistic structure of a request to the corresponding representations. The work is being done in the context of a prototype system that will allow physicians and other health professionals to access information in a computerized textbook of hepatitis through natural language dialogues. The interpretation of natural language queries is derived from DIAMOND/DIAGRAM, a linguistically motivated, domain-independent natural language interface developed at SRI. A text access component is being developed that uses representations of the propositional content of text passages and of the hierarchical structure of the text as a whole to retrieve relevant information.

CONTENTS

A. INTRODUCTION

Recent research in computer-based natural language interface systems has demonstrated the feasibility and value of their use for the retrieval of data in formatted files. Extending these capabilities to unformatted natural language texts can provide access to a broad new range of information. The simplest strategy would be to convert the relevant words in the request into key words and search for their occurrence using some sort of inverted file addressing strategy. However, this procedure, which is used extensively in full-text retrieval systems, is based exclusively on similarity in the form of words and not on their content. Our work is aimed at improving the sophistication of text access facilities by adapting techniques for linguistic analysis, representation of knowledge, inferencing, and matching developed in computational linguistics and artificial intelligence. Accordingly, we are developing a methodology for representing textual information and procedures for relating the linguistic structure of a request to the corresponding representations.

Our initial objective is the implementation of a prototype system that is intended to provide physicians and other health professionals with flexible and focused access, through natural language dialogue, to passages in a text containing information on hepatitis.* The system has two major components: (1) a body of text together with representations of both the propositional content of the text passages and their hierarchical relations; and (2) a natural language interface for analyzing a wide variety of English utterances and matching them with appropriate text passages. The decision to focus our research in this way reflects our judgment that a system is essential for evaluating our progress and that our efforts must be guided by specialists in the subject matter of the text.**

----

* The text we are using is the "Hepatitis Knowledge Base," developed at the National Library of Medicine as part of a program of research on transferring new research findings to medical practioners (Bernstein et al., 1980).

Dynamic text retrieval facilities of this kind are particularly appropriate for scientists and other professionals who are doing research or making decisions that require searching through documents. Such facilities also provide a context in which to gather data on how people formulate and test hypotheses. Dialogue interactions can reveal the progressive refinements made by an investigator as he simultaneously clarifies his questions and evaluates possible answers. The more natural the access mechanism is and the more sharply focused the results it provides, the more insight we can get into the processes of analysis and synthesis the investigator goes through.

In the rest of this introduction, we briefly review the limitations of current document retrieval systems, research on natural language interfaces and knowledge representation, and the background and organization of the Hepatitis Knowledge Base.

1. Document Retrieval Systems

Current document retrieval systems are both awkward to use and diffuse in the responses they provide. They may be grouped into two major classes, depending on whether searching is for reference retrieval and mediated by citation data and index terms or to identify a particular set of passages through full-text searching. Both kinds of systems require the user to master a relatively complex query language, which is sufficiently difficult that a group of professional information specialists has emerged to serve as intermediaries. In the reference retrieval situation, after a set of documents is identified, the user still has to locate them physically to determine whether they are actually relevant. The inclusion of an abstract in the computer file and the possibility of matching against the words it contains may refine the search and increase the likelihood of success, but the ultimate assessment must be made off-line (cf. Lancaster 1979).

** Medical guidance for the research is being provided by William S. Robinson, M.D., Professor of Medicine, Division of Infectious Diseases, Stanford University, a specialist in hepatitis; and by Charles H. Clanton, M.D., Research Associate at Stanford, a specialist in the application of computers to medicine.

In full-text searching, an approach used extensively by the legal profession (cf. Larson and Williams 1980), the user must specify the exact words to be matched in the text, although synonym facilities and techniques for truncation (removing prefixes and suffixes) are available. It is possible to require that a set of words must co-occur in the same sentence or paragraph, but one cannot prescribe how the words in the set relate to each other syntactically or semantically. It is difficult to anticipate how an author will describe a given topic; thus, many passages are likely to be missed. In addition, since the same words can be used in quite different contexts, many of the passages retrieved do not deal with the topic at all.

Current research in information science is addressing some of these problems (cf. Cooper 1978; McGill and Huitfeldt 1979; Doszkocs and Rapp 1979; Dattola 1979; O'Connor 1980), but recent developments in computational linguistics and artificial intelligence show more promise of providing new and sophisticated techniques for working with texts. In particular, in our research we are combining the use of natural language dialogue capabilities with new procedures for knowledge representation and access.

2. Natural Language Interfaces and Knowledge Representation

The emergence of practical natural language interfaces has been an important advance in computational linguistics (Harris 1977, 1979; Hendrix 1977a,b; Hendrix et al. 1978; Kolodner 1980; Petrick 1978; Templeton 1979; Thompson and Thompson 1978; Waltz 1978). They make it possible for a person to interact easily and directly with a data base via English questions, commands, and statements. However, many of the systems developed have tended to be ad hoc, application-specific, and difficult to extend. For working with text, it is essential to have a comprehensive grammar of English and facilities for handling the complex semantic relations entailed. DIAMOND/DIAGRAM, the system we are using in the current project, was developed at SRI for research on frontier issues in discourse--handling knowledge and belief, goals and plans in communication between people and computers (Robinson et al. 1980, Robinson 1980; Grosz 1981). However, it has proved to be robust enough for practical applications (Moore 1979a; Konolige 1979). DIAMOND/DIAGRAM is both linguistically motivated and domain-independent; moreover, it makes use of special semantic operators to provide semantic representations of the knowledge relevant for an application area.

The representation of knowledge in a computational system is one of the key problems addressed in artificial intelligence research. A great deal of work has been done on "assertional" representations, including work on production systems (Newell and Simon 1972; Davis and King 1976), semantic networks (Quillian 1968; Simmons 1973; Woods 1975; Hendrix 1979), and variants of first-order predicate calculus (Green 1969; Nash-Webber and Reiter 1977). Recently there have been advances in the representation of knowledge in specific areas, including belief states (Moore 1979b), temporal information (Lamport 1978), and other complex domains (McDermott and Doyle 1980; Weyhrauch 1980; Hayes 1979; Kolodner 1980). In addition, there has been progress in the representation of large bodies of highly structured knowledge (Minsky 1975; Bobrow and Winograd 1977; Schank and Abelson 1977; Schank et al. 1980), including process models (Robinson et al. 1980).

In our research we are using techniques derived from this body of AI research for encoding summary characterizations of the information in each passage of a full-text data base, as well as encoding the structure of the entire data base itself. This approach to representation provides a much more precise specification of the contents of a passage than is possible with a list of index terms. Matching and inferencing capabilities developed in artificial intelligence research (e.g., Fikes and Hendrix 1977) can then be used to retrieve these passages in response to a particular request.

3. The Hepatitis Knowledge Base

The Hepatitis Knowledge Base (HKB) synthesizes, in text form, the current medical information about viral hepatitis (Bernstein et al. 1980). It was developed as a prototype computer-based information transfer system to support a physician's decisions about diagnosis, prognosis, and treatment. The contents are validated by a panel of recognized experts in hepatitis, who review current literature and update it regularly. Consequently, it satisfies three prime requirements for use by specialists in the field: timeliness, comprehensiveness, and validity. In addition, the HKB is systematically organized in a hierarchical structure a fact of key importance in the approach to text retrieval outlined below.

B. PROCEDURES FOR TEXT ACCESS

Our text access procedures fall into five categories: representing the informational structure of the HKB text in a canonical language; linguistic analysis of a request and its translation into propositional form; resolving anaphoric expressions; reducing the request to primitives via synonymy and other inferential relations; and matching the request against the text structure. The first one and last two are the central issues of our research, so they will be described in the greatest detail, although we will discuss the problem of anaphoric reference briefly at the end.

The syntactic analysis and translation of a request into propositional form is done by SRI's DIAMOND/DIAGRAM system, which is described elsewhere (Robinson et al. 1980; Robinson 1980; Konolige 1979; for earlier work on the system see Paxton 1977; Walker 1978). It is worth pointing out the advantage of such a conversion over traditional index-term approaches to information retrieval. It can be illustrated by considering the two questions:

What indicates renal failure?
What does renal failure indicate?

An index-term search will scan these queries and, in both cases, come up with the keywords "indicate" and "renal failure." Thus the portions of text retrieved in each instance will be the same. Yet very straightforward syntactic processing will discover that in the first case renal failure is the Subject, in the second case the Object. In the first case, the proposition "indicate(?X,RENAL-FAILURE)" would be produced, in the second "indicate(RENAL-FAILURE,?X)". A passage stating that renal failure is a sign of fulminant viral hepatitis would have the proposition "indicate(RENAL-FAILURE,FULMINANT-VIRAL-HEPATITIS)" associated with it, instead of merely the keywords "renal-failure" and "indicate." It would thus be retrieved in the second example but not the first. The propositional approach thus allows much finer discrimination.

1. Representing the Text Structure

Two aspects of the informational structure of the HKB are represented: the information content of the individual passages, and the hierarchical organization of the whole. A summary of each passage is expressed as a proposition or a conjunction of propositions. For example, a paragraph describing a particular procedure for the prevention of hepatitis B could have associated with it the proposition "immunize(GAMMA-GLOBULIN,HEPATITIS-B)". Similarly, a passage concerned with the etiology of the disease could have the proposition "transmit(TRANSFUSION,HEPATITIS-B)". These expressions, together with representations of their hierarchical organization and pointers to their associated passages, constitute what we call the "text structure."

To make the matching process between the request and the text feasible, the text structure must be expressed in some canonical form. For this purpose we are doing a taxonomic analysis of the text, of the sort that Sager and her colleagues (cf. Sager 1978) have done in the domain of pharmacology, to determine the basic semantic word classes and the co-occurrence relations between the classes. For each set of synonyms in each class we pick one to be the canonical or primitive form. In this way, a language of primitives for the domain is built up. In addition, hyponymy, or inclusion, relations between terms are determined and represented. We then encode summaries of the selected portions of the HKB in propositional form, using the primitive or canonical terms as the predicates and arguments.

As an illustration, consider the following passage:

> The commonplace case of typical viral hepatitis is established by recognition of the inflammatory process of the liver (by its enlargement and tenderness) coupled with laboratory evidence of hepatic parenchymal damage (e.g., elevated aminotransferase activities in the serum) and some degree of impaired hepatic excretion of bilirubin (as indicated by changes in serum bilirubin and urine bilirubin and urobilinogen). In some there will be serologic evidence of a hepatitis virus (hepatitis B surface antigen or antibody).

This passage shows that such phrases as "inflammatory process of the liver," "enlargement (of the liver)," "tenderness (of the liver)," "hepatic parenchymal damage" and "elevated aminotransferase activities in the serum" can all occur as the subject of the verb "establish," whose object can be phrases like "typical viral hepatitis." Analysis of other passages reveals a large class of terms that can occur as subject of "establish" and suggests a characterization of the entire class, such as "evidence" or "symptoms and signs." Analysis shows that the object of "establish" can be one of several "forms." Furthermore it reveals a number of other words that can occur in the same environment as "establish," such as "show" and "indicate." The first stage in the analysis of the text is the determination of these word classes. The second stage is the determination of synonymy relations, and the choice of one of the terms as the canonical term. For example, since the verbs "establish," "show," and "indicate" function as synonyms in this domain, we have chosen one, "indicate," for use in representing the content of passages. The summaries of the passages are then expressed in terms of canonical predicates.

2. Reducing the Request to Canonical Form

The text structure is represented by means of a restricted language of canonical predicates and entities. The user's requests, by contrast, can be expressed in quite complex ways. If we are to match the request with the appropriate portion of the text structure, we must reduce the request to canonical form.

The simplest sort of reduction is when the terms the user employs are synonyms of the canonical predicates. In this case, the request is translated into the canonical synonym. For example, since a physician uses the term "reactivated" for "relapsing," we can treat the two as synonymous and the latter as canonical. The equivalence can be expressed by a rule in the system of the form:

(1) relapsing(x,y) <--> reactivated(x,y).

In addition to synonymy relations, the inferential techniques also handle hyponymy, or inclusion, relations. Frequently, a user asks a question in general terms when what he has in mind is quite specific. The system must then determine the specific interpretation. For example, the user may ask:

(2) How could I confirm that the hepatitis in this patient is a reactivated form of his previous disease?

There may be a number of ways to confirm something in the medical domain, depending on what needs to be confirmed. One way is diagnosis, and "diagnose" is a canonical predicate. There is a rule in the system of the form:

diagnose(x,y,p) --> confirm(x,p(y))

That is, if x diagnoses y as p, then x confirms that p is true of y. This rule would be applied for this request, so that the matching process seeks places in the text that discuss diagnostic tests that <u>diagnose</u> a form as relapsing rather than <u>confirm</u> it as relapsing. We may say that the proposition involving "diagnose" is an <u>interpretation</u> of the proposition involving "confirm." If there is more than one possible interpretation of a request, matches are sought for all of them.

Often the user conveys information in his request that the system cannot deal with. For example, in request (3):

(3) Give a brief summary of the types of hepatitis and their primary modes of transmission.

the physician asks about the "primary modes of transmission." The current system cannot make distinctions as subtle as the one between primary and secondary modes of transmission. Therefore, the property of the modes M, "primary(M)", is simply ignored.

A related problem results from the user embedding propositions to be matched inside higher predicates that are meaningless to the system. In request (2), the system seeks a match with the propositions involving the predicate "reactivated." However, in the initial representation of the request, this proposition is embedded in the operator "could." The relevant passages in the text structure are represented in terms of the predicate "relapsing," which is related to "reactivated," according to rule (1), but not in terms of the operator "could." Our approach to this problem is to ignore the embedding operator when it has no canonical translation and to seek a match only on the arguments. However, these ignored propositions are still available for the anaphoric reference resolution routines described below.

This approach is only as good as the rules that support it. If, for example, we have not anticipated that a physician might use "reactivated" for "relapsing" or "HAA" for "HBsAg", then we will be unable to interpret his request correctly.

3. The Matching Processes

In the simplest cases, the reduction processes reduce the request to canonical form, and that form directly matches one or more expressions in the text structure. For example, the reduction processes reduce request (4):

(4) Does positive HAA indicate a reinfection with hepatitis B?

to expression (5).

(5) indicate(occur(HB), infect(HBsAg)),

where HB is hepatitis B. Paragraph 239 in the HKB begins:

In man, as well as in experimentally infected chimpanzees, exposure to HBV results in the appearance of specific markers of HBV infection which have characteristic sequential relationships to time of exposure, the development of clinical manifestations, and recovery or the development of persistent infection. In the most common pattern HBsAg is the first detectable marker, appearing one to four or more weeks after exposure during the incubation period.

This passage thus has (5) in its summary. Therefore, matching the request with the text structure presents no difficulties.

Unfortunately, the matching process is not always so straightforward. Request (6) illustrates two serious problems--metonymy and the representation of refinements:

(6) Give a brief summary of the types of hepatitis and their primary modes of transmission.

In paragraph 2 of the textbook, the three types of hepatitis are defined and their characteristics are summarized. The summary of the paragraph is therefore

definition(TYPES,X), summary(TEXT,Y), characterize(Y,TYPES).

That is, the definition of the types is X, and the text summarizes the characteristics Y of the types. Analysis of the first conjunct of the request would have produced the expression

summary(?X, TYPES).

That is, the request is for a summary of the types, but the textbook contains only a summary of the <u>characteristics</u> of the types. Of course, as fluent speakers of English we recognize that is what the request is really for and that "summary of the types" is really an example of metonymy, or the omission of an intermediate function, "characteristics of."

This phenomenon is extremely prevalent (cf. Nunberg 1978), and an effective system must be able to deal with it. There are cases where such expressions are idiomatic and can be expanded in the linguistic processing. But for the general case, we employ an algorithm that is suggested by the work of Nunberg (1978) and Hobbs (1977): If we succeed in matching the predicate and all but one of the arguments of the request expression with a text structure expression, we search the properties of the remaining argument in the text structure expression for a proposition that has as one of its arguments the remaining argument in the request expression. For example, in processing request (6), we match on the predicate "summary," and the first arguments, "?X" and "TEXT," match (since the variable ?X will match with anything). Since "Y" and "TYPES" do not match, we look at the other properties of Y, in this case, "characterize(Y,TYPES)", and find that the unmatched argument, "TYPES", is one of its arguments. This process can be chained to interpret metonymic expressions with more than one omitted intermediate function.

The second problem illustrated by request (6) is in the second conjunct, which may be represented

summary(?X,M), mode(M,transmit(TYPES)) .

The information in paragraph 2 of the HKB, characterizing the types, is refined in the subsequent paragraphs. At the next level of refinement, we have paragraphs 7ff, which discuss laboratory tests that differentiate the types, and paragraphs 382 and 383, which discuss various manifestations that differentiate the types. Among these manifestations are prior infection, mode of transmission, and incubation period. Thus, we have a structure like the following, where the links represent refinement relations:

```
      summary(TEXT,Y), characterize(Y,TYPES)
                                /        \
                  _____________/          \
                 /                         \
differentiate(LAB-TESTS,TYPES)    differentiate(MANIFESTATIONS,TYPES)
               /  |  \                                /     |     \
                 ...                                /      |      \
                                            prior    modes    incubation
                                          infection    of       period
                                                  transmission
```

As before, we have a partial match on "summary", but not on the arguments. Here we must look at the refinements of the characteristics.

This example suggests the following method: If there is no match between the request expression and the text structure expression, follow the refinement links from the latter and test for a match again. Apply this recursively until a match is found. This method has been described in a top-down fashion, but in practice it proves more efficient to proceed bidirectionally, both downward from occurrences of "summary" in the text structure and upward from occurrences of "modes of transmission," looking for an intersection.

There are occasions when sophisticated matching techniques fail to achieve a complete match between request and text. In these cases, the various partial matches are rated and returned. For example, consider request (7):

(7) What is the chance that this relapsing form is the cause in this patient?

With an inference rule relating "chance" to "incidence," we derive the request expression

incidence(?X,cause(F,Y)), relapsing(F,D).

There are two paragraphs that may be relevant. Paragraph 1415 in the HKB discusses possible causes of relapsing forms:

cause(C,F), relapsing(F,D).

There is a match on "cause" and "relapsing," but the order of the arguments is wrong, so the partial match is not particularly good. Paragraphs 1420-2 discuss the incidence, or likelihood, of a relapsing form occurring:

incidence(FR,F), relapsing(F,D),

where FR is the frequency. Here we have a match on "incidence" and "relapsing" but not on "cause." The best we can do in this case is to score the two partial matches with a rating function, and return both passages. The exact scoring function is a matter of continuing research, but among the factors that need to be taken into account are how much reduction had to be done to convert the request into a canonical form, how complete the match was, how many intermediate functions were required in the case of metonymy, how many refinement links had to be followed, and, as described below, how highly focused the text structure expression was that was found.

4. Using the Hierarchical Structure of the Text

We are experimenting with three methods for taking advantage of the hierarchical structure of the text. The first is the one described above, using the refinement links that encode the hierarchy in the matching algorithm itself. In the second, the hierarchical structure of the text is used as a focus mechanism. Grosz (1977, 1979) has shown that the hierarchical structure of a task model can be used to good effect for focusing on the entities and task steps referred to in task-oriented dialogues. Similar methods can be employed in text retrieval, as illustrated by the following example.

Suppose a user asks:

> Is melena a sign of fulminant viral hepatitis?

Having obtained the text he wants, the user then asks

> How significant is abnormal prothrombin time?

Information on prothrombin time can be found in several places in the text, but the system assumes that the user is still interested in fulminant viral hepatitis, so those passages are presented first. There is generally a high degree of continuity in user requests. The focus mechanism provides a way to sharpen the system response when the request itself is not completely specified.

The third method for using the hierarchical structure of the text is what might be called "subsequent text access." Once the user has reached a part of the text he is interested in, he will have a number of other needs for moving around in the text. Among the facilities we provide initially are the following:

1. The user is able to locate his place with respect to the whole text by printing out the part, chapter, section, and subsection titles of the paragraph he is currently looking at. The user can also print out the table of contents one or more levels down for any of these regions, to find related passages.

2. The user can continue sequentially through the file, by requesting the next paragraph, the previous paragraph, and so on.

3. Where a number of possibly relevant passages have been found and rated by likely relevance, a user might, after looking at the most likely, want to look at the next most likely. Alternatively, he might want to add a further restriction to his original request, for example:

> Give me only passages involving drug addiction.

This initiates a more refined search through the possibilities already found, using variants of the techniques of initial text access to reduce their number and to reorder them in certain ways.

5. Resolution of Anaphoric Expressions

We are experimenting with several techniques for the resolution of anaphoric expressions. First consider the pronoun "their" in request (3).

(3) Give a brief summary of the types of hepatitis and their primary modes of transmission.

A fairly simple algorithm, described in Hobbs (1976) and implemented in the DIAMOND system, travels up the parse tree from "their" searching successive clauses and noun phrases for a plural noun phrase. Hobbs has shown that the algorithm is remarkably effective in general, and it works in this particular example, identifying the antecedent as "the types of hepatitis."

To resolve definite noun phrases, the representation of the previous utterances in the dialog is searched for an entity whose properties include the explicit properties of the entity referred to by the definite noun phrase. A simple example is provided by the phrase "this patient" in

(8) Could the hepatitis in this patient be some reactivated or chronic form of his prior disease?

Let us assume that in processing the context provided by the physician, an entity P was created for the patient, and, among his other attributes such as age and sex, it was recorded that he was a patient--"patient(P)". This proposition is located and "this patient" is assumed to be coreferential with P.

Suppose the next question asked by the physician is one we considered earlier:

(7) What is the chance that this relapsing form is the cause in this patient?

The phrase "this relapsing form" poses a more difficult problem. Request (8) referred to a "reactivated form." However, the reduction processes derive the fact that the reactivated form is also a relapsing form, since "relapsing" is a canonical predicate. Its properties thus include the explicit property contained in the phrase "this relapsing form," and the algorithm will find the correct antecedent.

The final example illustrates the power of this simple algorithm, provided the rest of the analysis of the dialogue is rich enough to support it. The phrase "this possibility" in request (9):

(9) What diagnostic tests should I use to confirm this possibility?

conveys only the information that it refers to something that is possible, an X such that "possible(X)". The word has not occurred previously in the dialogue, but the word "could" in request (8) is equivalent to "is it possible that." The reduction processes can derive from the representation of (8)--"could(H,reactivated(H,D))"--the equivalent expression "possible(reactivated(H,D))". The definite-noun-phrase algorithm then identifies the antecedent of "this possibility" as "that the hepatitis in this patient is a reactivated (or relapsing) form."

Where more than one candidate antecedent is found by this algorithm, we simply choose the most recently occurring candidate. It has been shown that this procedure is not always accurate (Charniak 1974), but in the absence of a deeper analysis of the structure of the interaction on which to base a preference (cf. Grosz 1977, Hobbs 1975, Lockman 1978), the algorithm gives reasonably good results.

Note that in text retrieval applications, it is frequently not crucial to resolve anaphoric expressions. They often refer to the specific case the physician is dealing with--e.g., "this patient"--whereas what is desired from the textbook is general information about cases of that class. The identification of relevant passages does not depend on how such anaphoric expressions are resolved.

C. SIGNIFICANCE

Our research constitutes an experiment in what might be called "computational pragmatics." One of the most important problems in natural language processing is relating an utterance in an ongoing dialogue to what is happening in the participants' environment, that is, determining the purpose of the utterance. Previously, this problem has been attacked most successfully in domains where significant portions of the environment can be modelled in a precise fashion in terms of the participants' goals and plans. In these applications, a request is matched, directly or through inferences, with a preexisting plan or a plan which is extended on the spot, in order to determine the intent of the request. In our research, the hierarchical organization of the text offers a similarly well-structured, although very different, model of a portion of the environment. Although static, it is much larger and less precise than other such models, and thus forces us to make important generalizations in the techniques of computational pragmatics.

D. REFERENCES

Bernstein, LM, Siegel, ER, and Goldstein, CM. The Hepatitis Knowledge Base. Annals of Internatl Medicine, 1980, 93, 165-222.

Bobrow, D and Winograd, T. An Overview of KRL, A Knowledge Representation Language. Cognitive Science, 1977, 1, 3-46.

Charniak, E. Toward a Model of Children's Story Comprehension. AI-TR-266, Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, Massachusetts, 1974.

Cooper, WS. Indexing Documents by Gedanken Experimentation. J. American Society for Information Science, 1978, 29, 107-119.

Dattola, RT. FIRST: Flexible Information Retrieval System for Text. J. American Society for Information Science, 1979, 30, 9-14.

Davis, R, and King, J. An Overview of Production Systems. In Machine Representation of Knowledge. Edited by EW Elcock and D Michie. Wiley, New York, 1976.

Doszkocs, TE, and Rapp, BA. Searching MEDLINE in English: A Prototype User Interface with Natural Language Query, Ranked Output, and Relevance Feedback. Proceedings of the ASIS Annual Meeting, Volume 16. Edited by RD Tally and RR Deultgen. Knowledge Industry Publications, New York, 1979. Pp. 131-139.

Fikes, RE and Hendrix, GG. A Network-Based Knowledge Representation and its Natural Deduction System. Proceedings of the Fifth International Joint Conference on Artificial Intelligence. Carnegie-Mellon University, Pittsburgh, Pennyslvania, 1977. Pp. 235-246.

Green, CC. The Application of Theorem Proving to Question Answering. Artificial Intelligence Project Memo AI-96, Stanford University, Stanford, California, 1969.

Grosz, BJ. The Representation and Use of Focus in a System for Understanding Dialogs. Proceedings of the Fifth International Joint Conference on Artificial Intelligence. Carnegie-Mellon University, Pittsburgh, Pennsylvania, 1977. Pp. 67-76.

Grosz, BJ. Focusing and Description in Natural Language Dialogues. In Proceedings of the Workshop on Computational Aspects of Linguistic Structure and Discourse Setting. Edited by AK Joshi, I Sag, and B Webber. Cambridge University Press, Cambridge, Massachusetts, 1981.

Harris, LR. User Oriented Data Base Query with the ROBOT Natural Language Query System. International Journal of Man Machine Systems, 1977, 9, 697-713.

Harris, LR. Experience with ROBOT in 12 Commercial Natural Language Data Base Query Applications. Sixth International Joint Conference on Artificial Intelligence. Stanford University, Stanford, California, 1979. Pp. 365-368.

Hayes, PJ. The Logic of Frames. In The Frame Reader. De Gruyter, Berlin, 1979.

Hendrix, GG. Human Engineering for Applied Natural Language Processing. Proceedings of the Fifth International Joint Conference on Artificial Intelligence. Carnegie-Mellon University, Pittsburgh, Pennyslvania, 1977. Pp. 183-191. (a)

Hendrix, GG. The LIFER Manual: A Guide to Building Practical Natural Language Interfaces. Technical Note 138, Artificial Intelligence Center, Stanford Research Institute, Menlo Park, California, February 1977. (b)

Hendrix, GG. Encoding Knowledge in Partitioned Networks. In Associative Networks - The Representation and Use of Knowledge in Computers. Edited by NV Findler. Academic Press, New York, 1979.

Hendrix, GG, Sacerdoti, ED, Sagalowicz, D, and Slocum, J. Developing a Natural Language Interface to Complex Data. ACM Transactions on Database Systems, 1978, 3, 105-147.

Hobbs, J. A General System for Semantic Analysis of English and Its Use in Drawing Maps from Directions. American Journal of Computational Linguistics, Microfiche 32, 1975.

Hobbs, J. Pronoun Resolution. Research Report 76-1, Department of Computer Sciences, City College, City University of New York, New York, New York, August 1976.

Hobbs, J. Coherence and Interpretation in English Texts. Proceedings of the Fifth International Joint Conference on Artificial Intelligence. Carnegie-Mellon University, Pittsburgh, Pennsylvania, 1977. Pp. 110-116.

Kolodner, JL. Organizing Memory and Keeping It Organized. Proceedings of the First Annual National Conference on Artificial Intelligence. American Association for Artificial Intelligence, Stanford University, Stanford, California, 1980.

Konolige, KG. A Framework for a Portable Natural Language Interface to Large Data Bases. Technical Note 197, Artificial Intelligence Center, SRI International, Menlo Park, California, October 1979.

Lamport, L. Times, Clocks, and the Ordering of Events in a Distributed System. Communications of the ACM, 1978, 21, 558-564.

Lancaster, FW. Information Retrieval Systems. Second Edition. Wiley, New York, 1979.

Larson, S, and Williams, ME. Computer Assisted Legal Research. In Annual Review of Information Science and Technology, Volume 15. Edited by ME Williams. Knowledge Industry Publications, White Plains, New York, 1980. Pp. 251-286.

Lockman, A. Contextual Reference Resolution in Natural Language Processing. Ph.D. Thesis, Department of Computer Science, Columbia University, New York, 1978.

McDermott, D, and Doyle, J. Non-Monotonic Logic I. Artificial Intelligence, 1980, 13, 41-72.

McGill, MJ, and Huitfeldt, J. Experimental Techniques of Information Retrieval. In Annual Review of Information Science and Technology, Volume 14, Edited by ME Williams. Knowledge Industry Publications, New York, 1979. Pp. 93-127.

Minsky, M. A Framework for Representing Knowledge. In The Psychology of Computer Vision. Edited by P Winston. McGraw-Hill, New York, 1975.

Moore, R. Mechanical Intelligence: Research and Applications. Final Technical Report, Artificial Intelligence Center, SRI International, Menlo Park, California, August 1979. (a)

Moore, R. Reasoning about Knowledge and Actions. Massachusetts Institute of Technology, Cambridge, Massachusetts. PhD Dissertation, 1979. (b)

Nash-Webber, B, and Reiter, R. Anaphora and logical form: on formal meaning representations for natural language. Proceedings of the Fifth International Joint Conference on Artificial Intelligence. Carnegie-Mellon University, Pittsburgh, Pennsylvania 1977. Pp. 121-131.

Newell, A, and Simon, HA. Human Problem Solving. Prentice-Hall, Englewood Cliffs, New Jersey, 1972.

Nunberg, G. The Pragmatics of Reference. Ph.D. Thesis, City University of New York, New York, 1978.

O'Connor, J. Answer Passage Retrieval by Text Searching. Journal of the American Society for Information Science, 1980, 31, 227-239.

Paxton, WP. A Framework for Speech Understanding. Technical Note 142, Artificial Intelligence Center, Stanford Research Institute, Menlo Park, California, 1977.

Petrick, SR. Automatic Syntactic and Semantic Analysis. In Proceedings of the Interdisciplinary Conference on Automated Text Processing (Bielefeld, German Federal Republic, 8-12 November 1978). Edited by J Petofi and S Allen. Reidel, Dordrecht, Holland, in press.

Quillian, MR. Semantic Memory. In Semantic Information Processing. Edited by M Minsky. MIT Press, Cambridge, Massachusetts, 1968. Pp. 227-270.

Robinson, JJ. DIAGRAM: A Grammar for Dialogues. SRI Technical Note 205, Artificial Intelligence Center, SRI International, Menlo Park, California. February 1980.

Robinson, AE, Appelt, DE, Grosz, BJ, Hendrix, GG, and Robinson, JJ. Interpreting Natural-Language Utterances in Dialog About Tasks. SRI Technical Note 210, Artificial Intelligence Center, SRI International, Menlo Park, California. March 1980.

Sager, N. Natural Language Information Formatting: The Automatic Conversion of Texts to a Structure Data Base. In Advances in Computers, Volume 17. Edited by MC Yovits and M Rubinoff. Academic Press, New York, 1978. Pp. 89-162.

Schank, R and Abelson, R. Scripts, Plans, Goals, and Understanding. Laurence Erlbaum Associates, Hillsdale, New Jersey, 1977.

Schank, RC, Lebowitz, M, and Birnbaum, L. An Integrated Understander. American Journal of Computational Linguistics, 1980, 6, 13-30.

Simmons, RA. Semantic Networks: Their Computation and Use for Understanding English. In Computer Models of Thought and Language. Edited by RC Schank and KM Colby. Freeman, San Francisco, 1973.

Templeton, M. EUFID: A Friendly and Flexible Front-end for Data Management Systems. Proceedings of the 17th Annual Meeting of the Association for Computational Linguistics. University of California at San Diego, La Jolla, California, 11-12 August, 1979. Pp. 91-93.

Thompson, BH, and Thompson, FB. Rapidly Extendable Natural Language. Proceedings of the 1978 Annual Meeting of the Association for Computing Machinery. ACM, New York, 1978. Pp. 173-182.

Walker, DE. (Ed.) Understanding Spoken Language. North-Holland, New York, 1978.

Waltz, DL. An English Language Question Answering System for a Large Relational Data Base. Communications of the ACM, 1978, 21, 526-539.

Weyhrauch, R. Prolegomena to a Theory of Formal Reasoning. Artificial Intelligence, 1980, 13, 133-170.

Woods, WA. What's In a Link? Foundations for Semantic Networks. In Representation and Understanding. Edited by DG Bobrow and A Collins. Academic Press, New York, 1975. Pp. 35-82.