

chological aspects, vision, robotics, theorem proving, applications, specialized systems, and game playing.

The deadline for submitting papers is May 1, 1980. Additional information can be obtained from:

Robert Balzer
AAAI Program Chairman
USC/Information Sciences Institute
4676 Admiralty Way
Marina del Rey, California 90291

Workshops on Semantics And Representation of Knowledge Being Held in New York

Semantics of natural language and methods for representing knowledge are the topics of a series of informal workshops held at the IBM Systems Research Institute in New York City. The workshops are open to anyone who is actively doing research on these topics in the fields of artificial intelligence, cognitive psychology, computational linguistics, data base design, linguistics, philosophy, psycholinguistics, and symbolic logic. The purpose of the workshops is to provide a forum for people who are working on these and related topics at universities and research institutions in New York and surrounding areas. Visitors from other locations are welcome to come, but we would like to establish a nucleus of participants who can attend regularly.

The format for the workshops is a meeting about once a month, normally on Friday mornings. The first meeting is scheduled for February 1, 1980 with a lecture, "Inferring Conceptual Meaning Structures," by Professor Sharon Salveter from SUNY at Stony Brook.

The IBM Systems Research Institute is providing a classroom and maintaining a computerized mailing list for the workshops. Anyone who is interested in attending and/or speaking at the workshops should send name, address, and a description of research interests to

John F. Sowa
IBM Systems Research Institute
205 East 42nd Street
New York, New York 10017
(212) 983-7231

Other Conferences

The *National Computer Conference*, sponsored by AFIPS, will be held in Anaheim, California, May 19-22, 1980. For further information, contact

Mr. Gerard Chiffrieller
AFIPS
1815 North Lynn St., Suite 800
Arlington, Virginia 22209

The *1980 LISP Conference*, sponsored by Stanford University, will be held at Stanford University in Palo Alto, California, August 24-27, 1980. For further information, contact

John R. Allen
Artificial Intelligence Laboratory
Stanford University
Stanford, California 94305

IFIP Congress '80, sponsored by IFIP, will be held in Kyoto, Japan, and Melbourne, Australia, October 6-17, 1980. For further information, contact

IFIP Congress '80
G.P.O. Box 880G
Melbourne, Victoria
AUSTRALIA 3001

ACM 80, this year's annual conference of the Association for Computing Machinery, will be held in Nashville, Tennessee, October 27-29, 1980. For further information, contact

Mr. Louis J. Fiora
ACM
1133 Avenue of the Americas
New York, New York 10036

Abstracts of Current* Literature

On the Difference Between Natural Language and High Level Query Languages

S. Jerrold Kaplan

Department of Computer and Information Science
Moore School of Electrical Engineering
University of Pennsylvania
Philadelphia, Pennsylvania 19104

ACM78, Proc. 1978 Annual Conf., Dec. 1978, 27-38.

Natural Language questions differ from most existing formal query languages in that they tend to admit a wider range of responses than their formal counterparts, and provide cues for selecting among the variety of appropriate responses. These differences reflect the fact that in Natural Language conversation, a respondent is expected to take an active role in the process of selecting and organizing responses, in contrast to formal query systems, where control of the interaction typically resides with the user or applications programmer. This paper explores some specific ways in which Natural Language questions are particularly well suited for their environment, and discusses the potential role of similar capabilities in formal query systems, particularly with respect to Natural Language Data Base query systems.

* Editor's note: The abstracts in this first issue are not as "current" as I would like, but, because we did not publish the *Journal* in 1979, there is a lot of material to get caught up on.

Impertinent Question-Answering Systems: Justification and Theory

Laurent Siklossy

Department of Information Engineering
University of Illinois at Chicago Circle
Chicago, Illinois 60680

ACM78, Proc. 1978 Annual Conf., Dec. 1978, 39-44.

A Question-Answering System (QAS) is impertinent if it tells a user something the user did not ask! Since a user rarely understands the full implications of a data base, especially a large one, his question is limited by his knowledge of the data base. Sometimes, a question similar to the one he originally asked could better satisfy his goals, but he did not ask it because he had no reason to do so. Hence, impertinence is a *necessary* feature of sophisticated QAS.

We describe a complete theory for impertinent QAS. The QAS is impertinent when, and only when, a small change in the question results in a large positive change in the answer, i.e., in the presence of a discontinuity. (The metrics are user dependent.) In impertinent QAS, the user's: "Why didn't you tell me?" answered by "BUT YOU DIDN'T ASK," no longer occur. The theory has been implemented in a sophisticated airline tariff system.

Natural Language Dialogue For Managing An On-Line Calendar

George E. Heidorn

Computer Sciences Department
IBM Thomas J. Watson Research Center
Yorktown Heights, New York 10598

ACM78, Proc. 1978 Annual Conf., Dec. 1978, 45-52.

This paper describes a project for studying the feasibility of developing systems which accomplish typical office tasks by means of human-like communication with the user. An actual dialogue with the initial version of a system that is being built for scheduling activities such as meetings is presented and then is used as a source of examples in explaining the operation of the system. The knowledge network is described, and the use of Augmented Phrase Structure Grammars for both the analysis and generation of English utterances in this system is discussed.

Some Comments on EQS, A Near Term Natural Language Data Base Query System

William A. Martin

MIT Laboratory for Computer Science
545 Technology Square
Cambridge, Massachusetts 02139

ACM78, Proc. 1978 Annual Conf., Dec. 1978, 156-164.

Problems and possibilities for near term natural language query systems are discussed, with emphasis on the author's own system, EQS. First, the general

objectives for near term systems in the areas of syntax, world knowledge, discourse, and problem solving are considered. Next, a comparison is made between the ATN parsing strategies in LADDER, ROBOT, PLANES, and EQS. Evidence for the importance of giving answers to queries not directly available in the data base is given together with some speculation on how the knowledge necessary for this might be acquired. After some general discussion of semantic data models, the semantic data model used in EQS is described, and the method used for building a semantic data model of an existing CODASYL data base is sketched. Lastly, the use of multiple levels of representation as a way to control complexity of such query systems is discussed.

The Robot System: Natural Language Processing Applied to Data Base Query

Larry R. Harris

Department of Mathematics
Dartmouth College
Hanover, New Hampshire 03755

ACM78, Proc. 1978 Annual Conf., Dec. 1978, 165-172.

In the early 1970's the natural language processing techniques developed within the field of artificial intelligence (AI) made important progress. Within certain restricted micro worlds of discourse it became possible to process a reasonable large class of English. These techniques have now been applied to the real micro world of data base query, allowing for information to be extracted from data bases by asking ordinary English questions. This paper discusses the importance or true natural language data base query and describes the ROBOT system, a high performance production level system already installed in several real world environments. The specific data structure requirements of the ROBOT system are discussed, as well as an extended type of data inversion that provides precisely the functionality required by the natural language parser.

Rapidly Extendable Natural Language

**Bozena Henisz Thompson and
Frederick B. Thompson**

Department of Information Science
California Institute of Technology
Pasadena, California 91109

ACM78, Proc. 1978 Annual Conf., Dec. 1978, 173-182.

A major thrust of artificial intelligence research is how to build knowledge of the application domain into computer systems. We investigate how the user himself can introduce his own expert knowledge into his data base system through rapid language extension so that it may then respond intelligently to his curt queries and commands. Illustrations of rapid language

extension using the REL System are presented and discussed.

An English Language Question Answering System for a Large Relational Database

David L Waltz

Coordinated Science Laboratory
University of Illinois
Urbana, Illinois 61801

Comm. of the ACM, 21, 7 (July 1978), 526-539.

By typing requests in English, casual users will be able to obtain explicit answers from a large relational database of aircraft flight and maintenance data using a system called PLANES. The design and implementation of this system is described and illustrated with detailed examples of the operation of system components and examples of overall system operation. The language processing portion of the system uses a number of augmented transition networks, each of which matches phrases with a specific meaning, along with *context registers* (history keepers) and *concept case frames*; these are used for judging meaningfulness of questions, generating dialogue for clarifying partially understood questions, and resolving ellipsis and pronoun reference problems. Other system components construct a formal query for the relational database, and optimize the order of searching relations. Methods are discussed for handling vague or complex questions and for providing browsing ability. Also included are discussions of important issues in programming natural language systems for limited domains, and the relationship of this system to others.

Developing a Natural Language Interface to Complex Data

**Gary G. Hendrix, Earl D. Sacerdoti,
Daniel Sagalowicz and Jonathan Slocum**

Artificial Intelligence Center
SRI International
Menlo Park, California 94025

ACM Trans. on Database Sys. 3, 2 (June 1978), 105-147.

Aspects of an intelligent interface that provides natural language access to a large body of data distributed over a computer network are described. The overall system architecture is presented showing how a user is buffered from the actual database management systems (DBMS's) by three layers of insulating components. These layers operate in series to convert natural language queries into calls to DBMS's at remote sites. Attention is then focused on the first of the insulating components, the natural language system. A pragmatic approach to language access that has proved useful for building interfaces to databases is described and illustrated by examples. Special language features that increase system usability, such as spelling correction, processing of incomplete inputs, and run-time

system personalization, are also discussed. The language system is contrasted with other work in applied natural language processing, and the system's limitations are analyzed.

Noun-Phrase Model and Natural Query Language

M. Sibuya, T. Fujisaki and Y. Takao

IBM Tokyo Scientific Center
1-11-32 Nagata-cho
Chiyoda-ku, Tokyo, JAPAN

IBM J. Res. Develop, 22, 5 (Sept. 1978), 533-540.

Basic considerations in designing a natural data base query language system are discussed. The notion of the noun-phrase data model is elaborated, and its role in making a query system suitable for general use is stressed. An experimental query system, Yachimata, embodying the concept, is described.

Interpretation of Natural Language in an Information System

H. Lehmann

IBM Scientific Center
Tiergartenstrasse 15
6900 Heidelberg, WEST GERMANY

IBM J. Res. Develop, 22, 5 (Sept. 1978), 560-572.

This paper discusses some of the linguistic problems encountered during the development of the User Specialty Languages (USL) system, an information system that accepts a subset of German or English as input for query, analysis, and updating of data. The system is regarded as a model for portions of natural language that are relevant to interactions with a data base. The model provides insight into the functioning of language and the linguistic behavior of users who must communicate with a machine in order to obtain information. The aim of application independence made it necessary to approach many problems from a different angle than in most comparable systems. Rather than a full treatment of the linguistic capacity of the system, details of phenomena such as time handling, coordination, quantification, and possessive pronouns are presented. The solutions that have been implemented are described, and open questions are pointed out.

The Meaning of "Of" and "Have" in the USL-System

Magdalena Zoeppritz

IBM Scientific Center
Tiergartenstrasse 15
6900 Heidelberg, WEST GERMANY

Technical Report 78.04.004, April 1978.

The paper shows how the transformational relationship between HAVE-sentences and OF-phrases is used to represent data contained in sentences with HAVE

as the main verb in the context of an information system using natural language to access a relational data base. An overview of the system first establishes the framework in which natural language processing is attempted. Then ways of representing HAVE are discussed with emphasis on the relation between HAVE and OF. The interpretation proposed and the interpretation process are illustrated by a list of representative queries and phrases against a small data base. In conclusion, this interpretation is extended to prepositional attributes with WITH and WITHOUT and problems are discussed.

Natural Language Access to a Melanoma Data Base

Martin N. Epstein

National Institutes of Health
Division of Computer Research and Technology
Bethesda, Maryland 20014

Donald E. Walker

Artificial Intelligence Center
SRI International
Menlo Park, California 94025

SRI Technical Note 171, Sept. 1978.

This paper describes ongoing research towards developing a system that will allow physicians personal access to patient medical data through natural language queries to support both patient management and clinical research. A prototype system has been implemented for a small data base on malignant melanoma. The physician can input queries in English that retrieve specified data for particular patients or for groups of patients satisfying certain characteristics, that perform simple calculations, that allow browsing through the data base, and that assist in identifying relations among attributes. The system supports dialogue interactions; that is, the user can follow a line of inquiry to test a particular hypothesis by entering a sequence of queries that depend on each other. Classes of questions that can be processed are described and examples using the system are given.

Information Structures in the Language of Science: Theory and Implementation

Naomi Sager and Lynette Hirschman

NYU Linguistic String Project
251 Mercer St.
New York, New York 10012

String Program Reports No. 12, Sept. 1978.

This report summarizes recent work on the relation of language to information structure, and on computer programs which convert the information in natural language documents in a given subject area into a structured data base. Operating on this data base, subsequent computer programs can retrieve specific facts and summarize the different types of information

present in the original documents. The technique is based upon the prior development of a natural language parsing system equipped with a comprehensive grammar of English. To this have been added procedures which reduce paraphrastically the variety of syntactic forms found in sentences and align parts of different sentences which have similar grammatical and informational standing in their respective sentences. When these procedures are applied to documents in the same subject area, the result is a table-like structure, called an information format, which contains the same information as the original documents, but arranged in a structured, rather than narrative, form.

Question Answering from Natural Language Medical Data Bases

Ralph Grishman and Lynette Hirschman

NYU Linguistic String Project
251 Mercer Street
New York, New York 10012

Artificial Intelligence 11 (1978), 25-43.

This paper describes a system for automatically answering questions about a collection of natural language medical records. The particular records used for an initial experiment were a set of 206 radiology reports. The implementation involves two major steps: manual determination of a suitable tabular structure (*information format*) for representing the information contained in the medical records, and automatic conversion of the natural language input (for either record or question) into a form corresponding to the data base. For the medical records the conversion into a data base is done by first performing a syntactic and transformational analysis of the sentences, followed by application of *formatting transformations*. The question-answering procedure has analogous initial steps but undergoes additional steps of processing to translate the question into a retrieval operation on the data base. Samples of the data base and of the question-answering procedure are shown.

Computers and Medical Language

M.G. Pacak and G.S. Dunham

Division of Computer Research and Technology
National Institutes of Health
Bethesda, Maryland 20205

Med. Inform. 4, 1 (1979), 13-27.

This review article presents the basic principles and problems of automated processing of medical language data. Work on automated processing of information in medical language is surveyed. References and a bibliography are provided as an introduction to the field.

Using a Formal Grammar in Human Factors Design of an Interactive Graphics System

Phyllis Reisner

IBM Research Laboratory
5600 Cottle Road
San Jose, California 95193

IBM Research Report RJ 2505, April 1979.

Formal grammatical description has not generally been applied in the human factors area, which traditionally draws on behavioral science for its methodology. This paper illustrates, by means of a detailed example, how formal grammatical description can be used in human factors.

The paper describes the human interface for two versions of an interactive graphics system. It then describes the "action language" of one version in terms of a production rule notation, and makes predictions about human performance based on the formal description of the first version and matching parts of the second. Some exploratory data is then given on results of testing some of the predictions. Since the predictions are based on general properties of the formal description, the technique should also be applicable to other "action languages."

Casting the Rendezvous Analyzer Rules into Augmented Transition Network Form

Timothy W. Finin

IBM Research Laboratory
5600 Cottle Road
San Jose, California 95193

IBM Research Report RJ 2146, Nov. 1978.

Rendezvous (Version 1) is a natural language query system for relational data bases. This document describes the conversion of the English language analyzer component of the RENDEZVOUS system from a phrase transformation rule form into an Augmented Transition Network form. The implementation (in APL) of a compiler for Augmented Transition Network grammars is also described.

Augmented Transition Network Grammars and Semantic Processing

Graeme D. Ritchie

Department of Computer Science
University of Edinburgh
James Clerk Maxwell Building
Mayfield Road
Edinburgh EH9 3JZ, SCOTLAND

Internal Report CSR-20-78, Jan. 1978.

The augmented transition network formalism is a special purpose programming language for organizing decisions during the analysis of an English sentence, which has been extensively used in syntactically based natural language programs for the past ten years. There have been a few suggestions recently that this

formalism could be used for semantically based programs. It seems unlikely that this proposal will be viable, since semantic categories are not the appropriate organizational units for an augmented transition network grammar.

Logic and Semantic Networks

Robert A. Kowalski

Department of Computing and Control
Imperial College of Science and Technology
London SW7 2BZ, ENGLAND

Amaryllis Deliyanni

Division of Electronics
University of Athens
Athens, GREECE

Comm. of the ACM 22, 3 (March 1979), 184-192.

An extended form of semantic network is defined, which can be regarded as a syntactic variant of the clausal form of logic. By virtue of its relationship with logic, the extended semantic network is provided with a precise semantics, inference rules, and a procedural interpretation. On the other hand, by regarding semantic networks as an abstract data structure for the representation of clauses, we provide a theorem-prover with a potentially useful indexing scheme and path-following strategy for guiding the search for a proof.

On Dealing with Quantification in Natural Language Utterances

Peter B. Sheridan

IBM Thomas J. Watson Research Center
P.O. Box 218
Yorktown Heights, New York 10598

Int. J. Man-Machine Studies 10 (1978), 367-394.

This paper presents some preliminary and still quite rudimentary ideas growing out of on-going studies into the problem of representation of quantified information in English, in terms of linked and nested "attribute-value" structures. Specifically, we indicate how such representation can be accomplished within the kind of attribute-value framework provided by Heidorn's NLP (Natural Language Processor) system.

A scheme for representing well-formed formulas of a sorted first-order logic (with descriptions) is first presented and serves as a bare conceptual framework for understanding the sequel. Then we indicate how certain commonly encountered English language quantificational utterances may be reduced to canonical forms representable within such an underlying scheme. A few of many still outstanding logico-linguistic issues (viz.: ambiguities connected with 'a', 'any' and 'some'; multiple predication; and anonymous quantification) are noted and discussed. Numerous examples are included.

Names And Naming: An Interdisciplinary Review

John M. Carroll

IBM Thomas J. Watson Research Center
P.O. Box 218
Yorktown Heights, New York 10598

IBM Research Report RC 7370, Oct. 1978.

The nature of names and their creation and use in naming is reviewed in the literature of three disciplines: philosophy, linguistics, and psychology. Convergent themes in this vast and yet inadequate literature suggest that a technical program of behavioral research may be feasible and productive in this area at the present time.

Process Models of Reference in Context

James A. Levin and Neil M. Goldman

USC/Information Sciences Institute
4676 Admiralty Way
Marina del Rey, California 90291

ISI Research Report 78-72, Oct. 1978.

Reference is a central issue for language comprehension and generation. After reviewing existing process models for comprehending and generating referring expressions, we present a general framework for context and reference processing. The context for reference processing is represented as a "Public Workspace." Reference processes access Public Workspace and modify its content, which is the set of concepts currently "on the table" as far as the current language interaction is concerned. Information from many different sources can be integrated in comprehending or generating referring expressions. Within this general framework, a new system for selectively generating referring phrases is developed. This system decides how much to express about a given concept in a given context.

Automatic Error Detection in Natural Language Words

R. Nussbaum and H.-J. Schek

IBM Scientific Center
Tiergartenstrasse 15
6900 Heidelberg, WEST GERMANY

Technical Report 78.06.005, June 1978.

An approach for the automatic error detection in natural language words without the use of a dictionary of permissible words is described. Automatically generated tables which describe permissible syllables and syllable sequences are used instead. These, on the other hand, are based on permissible initial and terminal consonant and vowel clusters. Therefore, an arbitrary string or erroneous word is tested whether it resembles a natural language word applying the above sequence tables. It is shown that under certain assumptions the error detection rate is approximately

80 by using relatively small tables with 70k bytes. The described principle may be applied recursively in order to obtain higher detection rates.

What Makes a Good Story?

Bertram C. Bruce

Bolt Beranek and Newman, Inc.
50 Moulton Street
Cambridge, Massachusetts 02138

Language Arts 55 (1978), 460-466.

This paper discusses two methods of text analysis used in research on children's understanding of stories, one based on story "grammars" and one on analysis of characters' plans. Using either one comes to similar basic conclusions. First, real stories have structures that can be identified and studied. Since it is a difficult task to recognize and use these structures in reading, children who have limited experience with real stories draw upon the reader's prior beliefs and expectations. The structure or connectivity of a story provides a framework for organizing appropriate prior beliefs. Third, the inherent complexity of story understanding, particularly the need to use prior beliefs in appropriate ways, means that there are usually several "correct" but different ways of understanding the same story. This suggests that children and adults may understand the same story in very different ways.

Conversational Coherency

Rachel Reichman

Bolt Beranek and Newman, Inc.
50 Moulton Street
Cambridge, Massachusetts 02138

CSR Technical Report No. 95, July 1978.

A major goal of this work is to specify some steps of the process by which participants maintain coherency in their conversations. The underlying element of the analysis is a construct called a 'context space'. Roughly, a group of utterances that refers to a single issue or episode forms the basis for a context space. Superficially, a conversation is a sequence of utterances; at a deeper level it is a structured entity whose utterances can be parsed into hierarchically related context spaces.

As a conversation proceeds, each conversant is building up a model of the conversation. Such a model includes the conversation's context space structure, a notion of the present discourse topic, and a list of items being focused upon. An important factor upon which the smoothness and coherency of a conversation depends is the lack of conflict between the respective models of the participants. Conflicts are prevented by a speaker following a number of high level syntactic and semantic relational rules that enable listeners to identify the speaker's discourse model. These rules are given in the paper, and, as the reader shall see, they

are highly integrated with the view of conversation presented here -- a hierarchy of related context spaces.

Elements of a Plan-Based Theory of Speech Acts

Philip R. Cohen and C. Raymond Perrault

Department of Computer Science
University of Toronto
Toronto, Ontario M5S 1A7 CANADA

Bolt Beranek and Newman Report 4057, Feb. 1979.

A theory of speech acts is developed that treats speech acts as operators in a planning system. The speech acts of requesting and informing are formally defined, and plans containing those speech acts are presented. Unlike previous (informal) speech act theories, this paper supplies justifications for the form and content of its speech act definitions. The adequacy criterion proposed here is that given an initial set of beliefs and goals, the speech act definitions should allow for the generation of plans for those speech acts that a person could issue appropriately under the same circumstances. This adequacy criterion is used in judging whether speech act definitions pass certain tests, in particular, the test of compositionality. For instance, since a speaker can request that a hearer do some arbitrary action, a speaker ought to be able to request that a hearer perform a speech act. A number of previous formulations of requesting and informing are shown to fail this test as applied to the planning of questions and to multi-party requests, such as "Ask Tom to open the door". In order to plan questions (requests to inform), two variations of informing speech acts are defined -- informing about the referent of a description and informing whether some proposition is true, which lead to WH and yes-no questions, respectively. It is argued that meta-theoretical principles should be developed that state how to formulate speech act definitions to pass the adequacy tests. One such principle is proposed and is shown to allow for the planning of multi-party requests.

Machine Translation and Machine-Aided Translation

W.J. Hutchkins

University of East Anglia
Norwich, UK

J. Doc. 34, 2 (June 1978), 119-159.

The recent report for the Commission of the European Communities on current multilingual activities in the field of scientific and technical information and the 1977 conference on the same theme both included substantial sections on operational and experimental machine translation systems, and in its *Plan of Action* the Commission announced its intention to introduce an operational machine translation system into its departments and to support research projects on machine

translation. This revival of interest in machine translation may well have surprised many who have tended in recent years to dismiss it as one of the "great failures" of scientific research. What has changed? What grounds are there now for optimism about machine translation? Or is it still a "utopian dream"? The aim of this review is to give a general picture of present activities which may help readers to reach their own conclusions. After a sketch of the historical background and general aims (Section I), it describes operational and experimental machine translation systems of recent years (Section II), it continues with descriptions of interactive (man-machine) systems and machine-assisted translation (Section III), (and it concludes with a general survey of present problems and future possibilities Section IV).

Knowledge-Based Machine Translation

Jaime Carbonell

Computer Science Department
Carnegie-Mellon University
Schenley Park
Pittsburgh, Pennsylvania 15213

Richard E. Cullingford

Electrical Engineering and Computer Science Dept.
The University of Connecticut
Storrs, Connecticut 06288

Anatole V. Gershman

Artificial Intelligence Project
Department of Computer Science
Yale University
New Haven, Connecticut 06520

Yale Computer Science Report 146, Dec. 1978.

This paper considers the possibilities for knowledge-based automatic text translation in light of recent advances in Artificial Intelligence. It is argued that competent translation requires some reasonable depth of understanding of the source text, and in particular, access to detailed contextual information. The following Machine Translation paradigm is proposed: First, the source text is analyzed and mapped into a language-free conceptual representation. Then, inference mechanisms apply contextual world knowledge to augment the representation by adding all information that was implicit in the input text. Finally, a natural-language generator maps the most relevant sections of the language-free representation into the target language. We illustrate this process with examples of English-to-Spanish and English-to-Russian translations, including some produced by SAM, a working system of computer programs at Yale University.

Memory Organization for Natural Language Data-Base Inquiry

Janet L. Kolodner

Artificial Intelligence Project
Department of Computer Science
Yale University
New Haven, Connecticut 06520

Research Report 142, Sept. 1978.

The CYRUS system is a natural language data-base query system containing biographical information about Cyrus Vance, Secretary of State of the United States. It also contains additional information associated with Vance that may be needed in order to answer questions about him, such as current events and limited biographical information about people he has been in close contact with. Questions input in English are answered by accessing conceptual structures containing this information.

This paper addresses the problems of knowledge organization and retrieval of information in the data-base. The conceptual structures used to represent, organize, and retrieve information are described, as well as the procedures used to answer natural language questions. Specific event information about the persons represented in the data-base must be stored; so must world events, general world knowledge, and rules about inference and reference resolution. The first half of the paper describes and explains some of the structures used for organizing this knowledge; the second half shows how CYRUS makes use of these structures.

Inference in the Conceptual Dependency Paradigm: A Personal History

Roger C. Schank

Yale University
Department of Computer Science
New Haven, Connecticut 06520

Research Report 141, Sept. 1978.

During a summer workshop (funded by the Sloan Foundation) that we ran at Yale, I tried to present some of the views that we hold with respect to the problems of representation of meaning, the making of inferences, and the function of higher level descriptions of the structure of knowledge, to an audience primarily consisting of social and cognitive psychologists. Most of the participants in the workshop were interested in our ideas on this subject. However, their background really had not prepared them to understand why we did what we did or how we came to do it. Consequently, I attempted to give them that background by retracing the steps in our research of the last ten years. I explained how we came to hold our current views on various subjects by showing what our initial assumptions were and how one position naturally led us to the next. Since most of the participants in

the workshop felt that these lectures provided the context that was necessary to help them to understand our current research, I felt that it would be of use to prepare a paper based on the lecture notes that I used in the workshop.

Interestingness: Controlling Inferences

Roger C. Schank

Yale University
Department of Computer Science
New Haven, Connecticut 06520

Research Report 145, Oct. 1978.

The problem of controlling inference is one of the most serious in Artificial Intelligence. New types of goal and plan inferences seriously compound the problem. This paper attempts to outline one possible solution to controlling inferences, namely following what is interesting and ignoring what is not.

Script Application: Computer Understanding of Newspaper Stories

Richard Edward Cullingford

Department of Computer Science
Yale University
New Haven, Connecticut 06520

Research Report 116, Jan. 1978.

This thesis describes a computer story understander which applies knowledge of the world to comprehend what it reads. The system, called SAM, reads newspaper articles from a variety of domains, then demonstrates its understanding by summarizing or paraphrasing the text, or answering questions about it. Since the knowledge structures SAM works with are conceptual and language-free, we have been able to add a limited machine-translation capability to SAM, as well.

SAM's knowledge of the world is encoded through the use of a representational construct called a *Script*. Scripts describe the stereotyped activities characteristic of socially ritualized situations such as going to stores, museums and restaurants, taking business trips and vacations, and attending banquets and birthday parties. SAM consults its Scripts to recognize the events a particular text refers to, to identify the participants in these events, and to fill in other events, not explicitly mentioned by a story, which can be plausibly inferred to have happened. In this process, it moves Scripts in and out of active memory on the basis of predictions it makes about what may be seen next.

Sam represents an attempt to build a complete, working story understander exploring an important source of knowledge about the world to find the connections which make a text "coherent," and to build a memory representation for the text from which natural-language outputs can be generated which indicate a reasonable depth of understanding. Since

Scripts model a knowledge structure that people constantly apply, both to cope with the world and to understand what they read, SAM embodies a theory of *context* and how context is to be used in the process of understanding.

Understanding Goal-Based Stories

Robert Wilensky

Dept. of Electrical Engineering and Computer Science
University of California
Berkeley, California 94720

Yale Computer Science Report 140, Sept. 1978.

Stories usually describe situations involving purposeful interactions between people. To understand stories, a reader needs a large body of knowledge about the kinds of situations that may occur, and must be able to use this knowledge to infer implicit aspects of a situation described by a text. Building a computer program that understands stories therefore requires an explicit description of this knowledge, and the construction of an inference mechanism with which to apply it.

This thesis is about understanding story situations involving goals. Most of the thesis is concerned with describing the categories of story situations that goals arise in, and with giving algorithms for detecting and processing the situations in each category. The most interesting categories of story situations involve multiple goals. Understanding a situation with a number of goals in it requires a reader to determine the interrelationships between those goals, and to infer the consequences of these relationships. Goal relationships form the basis for most of the situational categories presented here.

Knowledge of story situations and how they can be processed was used to build a story understanding program called PAM (Plan Applier Mechanism). PAM reads stories in English, and produces story representations that include the inferences needed to connect each story's events. To demonstrate that it has understood a story, PAM answers questions about the story and expresses the story from several points of view.

Subjective Understanding: Computer Models of Belief Systems

Jaime Guillermo Carbonell

Department of Computer Science
Carnegie-Mellon University
Schenley Park
Pittsburgh, Pennsylvania 15213

Yale Computer Science Report 150, Jan. 1979.

Modeling human understanding of natural language requires a model of the processes underlying human thought. No two people think exactly alike; different people subscribe to different beliefs and are motivated by different goals in their activities. A theory of subjective understanding has been proposed to account for subjectively-motivated human thinking ranging from ideological belief to human discourse and personality traits. A process-model embodying this theory has been implemented in a computer system, POLITICS. POLITICS demonstrates its understanding by answering questions in natural language question-answer dialogs.

POLITICS represents political ideologies as *goal trees* that structure the motivations of political actors and the relative importance of the various motivations. General purpose *counterplanning strategies* and more specific situational inference rules model the reasoning processes. Counterplanning strategies determine when, how and why political actors attempt various plans to thwart each other's goals, nullify each other's thwarting attempt, or end mutual hostilities. The subjective goal trees determine the focus of attention for the reasoning process. Only inferences that are relevant to the interests, goals and beliefs of the understander are pursued. This process drastically reduces the inference search space.

In addition to POLITICS, the theory of subjective understanding has been embodied in other process-models, also implemented as computer programs. These programs are TRIAD, a system that models human understanding of social conflict situations, and MICS, a system that models a participant in a natural language conversation. TRIAD generalizes the ideological conflict understanding of POLITICS to encompass all types of situations, ranging from domestic quarrels to judicial decisions and international disagreements. MICS uses goal trees to represent the conversational motivations of a person, and thus guides the course of the conversation towards the topics it is interested in discussing. All three systems presented in this dissertation are integrated understanding systems, incorporating subjective beliefs, reasoning strategies, and natural language analysis into a single, unified process.