

The *Third Annual Meeting of the Cognitive Science Society* will be held in Berkeley, California, August 19-21, 1981. Information about the meeting can be obtained from:

Robert Wilensky
Computer Science Division
Department of EECS
University of California, Berkeley
Berkeley, California 94720

The *Seventh International Joint Conference on Artificial Intelligence* will be held in Vancouver, British Columbia, Canada, August 24-28, 1981. [See *AJCL* 6:3-4, pg. 194.]

Questions about the technical program should be addressed to:

Roger C. Schank
Program Chairman, IJCAI-81
Department of Computer Science
Box 2158 Yale Station
Yale University
New Haven, Connecticut 06520
(203) 436-0606

General questions about the conference may be addressed to:

Pat Hayes
General Chairman, IJCAI-81
Department of Computer Science
Mathematical Sciences Building
University of Rochester
Rochester, New York 14627

The *1981 Annual Conference of the British Society for the Philosophy of Science* will be held at Wivenhoe House at the University of Essex, September 18-20, 1981. [See *AJCL* 6:3-4, pg. 194.] Further details of the conference may be obtained from:

Liaison Officer
University of Essex
Wivenhoe Park
Colchester CO4 3SQ, Essex, ENGLAND

The *1981 ACM Annual Conference* will be held at the Bonaventure Hotel in Los Angeles, California, November 9-11, 1981. [See *AJCL* 6:3-4, pg. 194.] For further information contact:

Mrs. A.C. Toni Shelter
Xerox Corporation, A3-49
701 South Aviation Boulevard
El Segundo, California 90245
(213) 679-4511 x1968

Abstracts of Current Literature

The Transformational Question Answering (TQA) System: Description, Operating Experience, and Implications

Fred J. Damerau
Mathematical Sciences Department
IBM Thomas J. Watson Research Center
Yorktown Heights, New York 10598
Research Report RC 8287, May 1980.

This paper sketches the structure of the TQA system with some examples, discusses some of the results of running the system in a user site during the year 1978, and outlines a few of the implications of a natural language query system for larger and more diverse data bases. It is apparent that a number of problems remain to be solved, but it also seems likely that a useful system can be constructed for some domains.

Design for an Intelligent Office System

Janet L. Kolodner
Department of Computer Science
Yale University
New Haven, Connecticut 06520
IBM Research Report RC 8385, July 1980.

This paper addresses the problem of designing an intelligent office system to help principals (managerial, administrative, or professional level office workers) with their work. A system is described which would use normative models of semi-routine office procedures to anticipate the user's next steps, to provide him automatically with files, documents, and forms he needs to complete a task, and to perform automated subtasks at the correct times. A number of intelligent functions that would perform some of the tasks secretaries generally do (scheduling, reminding, filing) are described, along with the knowledge necessary for performing those functions. In addition, an approach to building a more intelligent system which would understand natural language commands given by the user and understand and deal with routine mail is presented.

Memory Organization and Search Processes for Narratives

Michael G. Dyer and Wendy G. Lehnert
Department of Computer Science
Yale University
New Haven, Connecticut 06520
Research Report 175, April 1980.

BORIS represents the first system to integrate the knowledge-based inference techniques of scripts, plans, goals, and themes, within a single narrative understanding program. This report discusses techniques used by BORIS for memory representation and memo-

ry retrieval. Emphasis is placed on human question answering abilities, and the heuristics needed to simulate these phenomena. An example narrative processed by BORIS is discussed in detail and used to illustrate design decisions.

What Does it Mean to Understand Language?

Terry Winograd

Computer Science Department
Stanford University
Stanford, California 94305

Cognitive Science 4, 3 (July-Sept. 1980), 209-241.

In its earliest drafts, this paper was a structured argument, presenting a comprehensive view of cognitive science, criticizing prevailing approaches to the study of language and thought and advocating a new way of looking at things. Although I strongly believed in the approach it outlined, somehow it didn't have the convincingness on paper that it had in my own reflection. After some discouraging attempts at reorganization and rewriting, I realized that there was a mismatch between the nature of what I wanted to say and the form in which I was trying to communicate.

The understanding on which it was based does not have the form of a carefully structured framework into which all of cognitive science can be placed. It is more an orientation — a way of approaching the phenomena — that has grown out of many different experiences and influences and that bears the marks of its history. I found myself wanting to describe a path rather than justify its destination, finding that in the flow, the ideas came across more clearly. Since this collection was envisioned as a panorama of contrasting individual views, I have taken the liberty of making this chapter explicitly personal and describing the evolution of my own understanding.

My interests have centered around natural language. I have been engaged in the design of computer programs that in some sense could be said to "understand language," and this has led to looking at many aspects of the problems, including theories of meaning, representation formalisms, and the design and construction of complex computer systems. There has been a continuous evolution in my understanding of just what it means to say that a person or computer "understands," and this story can be read as recounting that evolution. It is long, because it is still too early to look back and say "what I was *really* getting at for all those years was the one basic idea that ..." I am too close and too involved in its continuation to see beyond the twists and turns. The last sections of the paper describe a viewpoint that differs in significant ways from most current approaches, and that offers new possibilities for a deeper understanding of language and a grasp on some previously intractable or

unrecognized problems. I hope that it will give some sense of where the path is headed.

Language and Memory

Roger C. Schank

Department of Computer Science
Yale University
New Haven, Connecticut 06520

Cognitive Science 4, 3 (July-Sept. 1980), 243-284.

This paper outlines some of the issues and basic philosophy that have guided my work and that of my students in the last ten years. It describes the progression of conceptual representational theories developed during that time, as well as some of the research models built to implement those theories. The paper concludes with a discussion of my most recent work in the area of modelling memory. It presents a theory of MOPs (Memory Organization Packets), which serve as both processors and organizers of information in memory. This enables effective categorization of experiences in episodic memory, which in turn enables better predictive understanding of new experiences.

"You Can't Miss It!": Judging the Clarity of Directions

Christopher K. Riesbeck

Department of Computer Science
Yale University
New Haven, Connecticut 06520

Cognitive Science 4, 3 (July-Sept. 1980), 285-303.

Recently there has been some very interesting AI work done on the representation of knowledge of large scale maps and their use in the construction of routes through some space (Kuipers, 1978; McDermott, 1980). The stress in this work has been on how spatial information is best represented and how reasoning is performed with the assumed representation. In this paper I would like to address the natural language processing of texts giving directions and make a claim that *during a casual first reading*, there actually does not need to be much spatial reasoning going on at all. When I read a written set of directions, I am primarily interested whether the directions seem clear and sensible - not in constructing a map or program specifying all the turns, distances, and locates that will be involved. Certainly there are times when I have to make some kind of route or map structure, such as when I am the subject in a spatial reasoning experiment, or when I am lost and ask someone on the street for directions. But I do not need to work so hard when I am given a piece of paper with a set of directions, and I know that I will have this piece of paper with me when I actually make the trip. I know that I can always read the directions again to figure out what to do, when I actually set out on the trip.

On Throwing Out the Baby with the Bathwater: A Reply to Black and Wilensky's Evaluation of Story Grammars

Jean M. Mandler

Department of Psychology
University of California, San Diego
La Jolla, California 92093

Nancy S. Johnson

State University of New York at Buffalo
Amherst, New York 14226

Cognitive Science 4, 3 (July-Sept. 1980), 305-312.

A number of criticisms of a recent paper by Black and Wilensky (1979) are made. (1) In attempting to assess the observational adequacy of story grammars, they state that a context-free grammar cannot handle discontinuous elements; however, they do not show that such elements occur in the domain to which the grammars apply. Further, they do not present adequate evidence for their claim that there are acceptable stories not accounted for by existing grammars and that the grammars will accept non-stories such as procedures. (2) They state that it has been proven that under natural conditions children cannot learn transformational grammars, which is a misrepresentation of the learnability proofs which have been offered. (3) Most important, they take an unduly narrow approach to story understanding by claiming that people only understand story content and do not have knowledge of story structure which is useful in comprehension or memory. Counter-evidence from the literature is cited which indicates that such knowledge is both useful and used, and a number of methods for assessing the psychological adequacy of structural models are discussed.

On Evaluating Story Grammars

David E. Rumelhart

University of California, San Diego
La Jolla, California 92093

Cognitive Science 4, 3 (July-Sept. 1980), 313-316.

In their recent article entitled "An Evaluation of Story Grammars," Black and Wilensky (1979) offer a critique of the recent work on this topic. They argue that story grammars (or story schemata as I prefer to call them) are not a productive approach to the study of story understanding, and they offer three main lines of argumentation. First, they argue that story grammars are not *formally* adequate in as much as most of them are represented as a set of context free rewrite rules which are known to be inadequate even for sentence grammars. Second, they argue that story grammars are not *empirically* adequate in as much as there are stories which do not seem to follow story grammars and there are nonstories which do. Finally, they argue that story grammars could not form an adequate basis for a comprehension model since in order to apply the grammar you need to have interpreted the

story. These arguments are in my opinion, indicative of a misunderstanding of the enterprise that I and others working on these issues have been engaged in. I believe that they are all based on a misunderstanding about what grammars might be good for and about how comprehension might occur. In this response, I wish to clarify the nature of story schemata as I understand them, clarify the nature of Black and Wilensky's misunderstandings and show how each of their arguments fails to address the important issues about story grammars and story schemata.

Natural Language Understanding

Avron Barr

Department of Computer Science
Stanford University
Palo Alto, California 94305

AI Magazine 1, 1 (Spring 1980), 5-10.

This is an excerpt from the *Handbook of Artificial Intelligence*, a compendium of hundreds of articles about AI ideas, techniques, and programs being prepared at Stanford University by AI researchers and students from across the country. This article, which is from the chapter on Natural Language Understanding, presents a brief sketch of the history of natural language processing research and gives an idea of the current state of the art. The other articles in the NL chapter of the Handbook include a historical sketch of machine translation, technical articles on grammars and parsing techniques, and an article on text generation. Finally, there are several articles describing the NL programs themselves. The Handbook also includes chapters on speech understanding and knowledge representation.

Utterance and Objective: Issues in Natural Language Communication

Barbara J. Grosz

Artificial Intelligence Center
SRI International
Menlo Park, California 94025

AI Magazine 1, 1 (Spring 1980), 11-20.

Two premises, reflected in the title, underlie the perspective from which I will consider research in natural language processing in this paper. First, progress on building computer systems that process natural languages in any meaningful sense (i.e., systems that interact reasonably with people in natural language) requires considering language as part of a larger communicative situation. In this larger situation, the participants in a conversation and their states of mind are as important to the interpretation of an utterance as the linguistic expressions from which it is formed. A central concern when language is considered as communication is its function in building and using shared models of the world.

Second, as the phrase "utterance and objective" suggests, regarding language as communication requires consideration of what is said literally, what is intended, and the relationship between the two. Recently, the emphasis in research in natural language processing has begun to shift from an analysis of utterances as isolated linguistic phenomena to a consideration of how people use utterances to achieve certain objectives. But, in considering objectives, it is important not to ignore the utterances themselves. A consideration of a speaker's underlying goals and motivations is critical, but so is an analysis of the particular way in which that speaker expresses his thoughts.

This paper examines three consequences of these claims for the development of language processing theories and the construction of language processing programs: (1) language processing requires a combination of language-specific mechanisms and general common-sense reasoning mechanisms, (2) language systems must be able to represent the beliefs and knowledge of multiple individual agents, and (3) utterances must be viewed as having effects along multiple dimensions.

Conversation as Planned Behavior

Jerry R. Hobbs and David A. Evans

Artificial Intelligence Center
SRI International
333 Ravenswood Avenue
Menlo Park, California 94025

Technical Note 203, Dec. 1979.

In this paper, planning models developed in artificial intelligence are applied to the kind of planning that must be carried out by participants in a conversation. A planning mechanism is defined, and a short fragment of a free-flowing videotaped conversation is described. The bulk of the paper is then devoted to an attempt to understand the conversation in terms of the planning mechanism. This microanalysis suggests ways in which the planning mechanisms must be augmented, and reveals several important conversational phenomena that deserve further investigation.

Metaphor, Metaphor Schemata and Selective Inferencing

Jerry R. Hobbs

Artificial Intelligence Center
SRI International
333 Ravenswood Avenue
Menlo Park, California 94025

Technical Note 204, Dec. 1979.

The importance of spatial and other metaphors is demonstrated. An approach to handling metaphor in a computational framework is described, based on the idea of selective inferencing. Three examples of metaphors are examined in detail in this light — a simple

metaphor, a spatial metaphor schema, and a novel metaphor. Finally, there is a discussion, from this perspective, of the analogical processes that underlie metaphor in this approach and what the approach says about several classical questions about metaphor.

Representation of Task-Specific Knowledge in a Gracefully Interacting User Interface

Eugene Ball and Phil Hayes

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

Technical Report CMU-CS-80-123, April 1980.

Command interfaces to current interactive systems often appear inflexible and unfriendly to casual and expert users alike. We are constructing an interface that will behave more cooperatively (by correcting spelling and grammatical errors, asking the user to resolve ambiguities in subparts of commands, etc.). Given that present-day interfaces often absorb a major portion of implementation effort, such a gracefully interacting interface can only be practical if it is independent of the specific tool or functional subsystem with which it is used.

Our interface is tool-independent in the sense that all its information about a particular tool is expressed in a declarative tool description. This tool description contains schemas for each operation that the tool can perform, and for each kind of object known to the system. The operation schemas describe the relevant parameters, their types and defaults, and the object schemas give corresponding structural descriptions in terms of defining and derived subcomponents. The schemas also include input syntax, display formats, and explanatory text. We discuss how these schemas can be used by the tool-independent interface to provide a graceful interface to the tool they describe.

Flexible Parsing

Phil Hayes and George Mouradian

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

Technical Report CMU-CS-80-122, May 1980.

When people use natural language in natural settings, they often use it ungrammatically, missing out or repeating words, breaking-off and restarting, speaking in fragments, etc. Their human listeners are usually able to cope with these deviations with little difficulty. If a computer system wishes to accept natural language input from its users on a routine basis, it must display a similar indifference. In this paper, we outline a set of parsing flexibilities that such a system should provide. We go on to describe FlexP, a bottom-up

pattern-matching parser that we have designed and implemented to provide these flexibilities for restricted natural language input to a limited-domain computer system.

A Parsing System for Montague Grammar with Lexical Transformations

Steven E. Tolkin

Computer and Communication Sciences
2076 Frieze Building
University of Michigan
Ann Arbor, Michigan 48109

Computer Studies in Formal Ling. N-24, Sept. 1980.

David Dowty's proposal to extend the grammar of Montague's *The proper treatment of quantification in ordinary English* by lexical transformations of verbs was implemented by modifying an existing system. Fundamental changes were made to the structure of the lexicon and to the ATN representation of the grammar. The viability of this implementation supports Dowty's category-changing approach to lexically governed transformations, and provides a base upon which further extensions to the grammar may be tested.

The ATN and the Sausage Machine: Which One is Baloney?

Eric Wanner

University of Sussex
Brighton BN1 9QL, ENGLAND

Cognition 8, 2 (June 1980), 209-225.

In a recent issue of *Cognition*, Lyn Frazier and Janet Dean Fodor proposed a new two-stage parsing model, dubbed the Sausage Machine (Frazier and Fodor, 1978). One of the major results which Frazier and Fodor bring forward in support of their proposal concerns a parsing strategy which, following Kimball (1973), they call Right Association. The center-piece of their argument concerns an interaction between this parsing strategy and another one, which they call Minimal Attachment. Frazier and Fodor (henceforth FF) provide interesting evidence that the language user makes tacit use of both strategies to resolve temporary syntactic ambiguities that arise during parsing. FF then proceed to argue that the existence of these strategies, as well as the apparent interaction between them, can be fully explained if we assume that the language user's parsing system is configured along the lines of the Sausage Machine. In FF's view, the Augmented Transition Network (ATN) runs a very poor second to the Sausage Machine, for according to FF's argument, it is impossible even to describe the two parsing strategies within the ATN framework. In effect then, FF are claiming that the Sausage Machine achieves explanation adequacy in this case while the ATN fails to reach the level of descriptive adequacy.

These are strong and potentially important claims. If correct, they obviously provide grounds for pursuing parsing models built along the lines of the Sausage Machine rather than the ATN. However, when FF's arguments are examined at close range, the comparison between parsing systems comes out rather differently than they claim. In particular, it appears that the Sausage Machine explanation of Right Association and its interaction with Minimal Attachment is empirically incorrect. The inadequacy of this explanation completely cancels the Sausage Machine's ability to describe the interaction between strategies that FF have observed. This follows because FF aspire to an explanation that renders independent description of the parsing strategies unnecessary. The Sausage Machine contains no apparatus for describing strategies. Hence, the failure to achieve explanatory adequacy automatically entails descriptive failure as well. In contrast, and in contradiction of FF's negative claim, the ATN can provide a perfectly general description for each strategy in terms of scheduling principles that constrain the order in which arcs in an ATN grammar are attempted. Moreover, when these scheduling principles are coupled with an ATN version of the grammar FF tacitly employed to generate their pivotal cases, FF's observations about the interactions between strategies are completely accounted for. Thus, although the ATN framework does not provide an explanation for either parsing strategy, it appears to achieve descriptive adequacy. Moreover, the descriptive framework of the ATN makes it possible to discern just what phenomena require explanation and to speculate in a reasonable way about the explanatory principles that underlie the parsing strategies FF have discovered.

An Improved Context-Free Recognizer

**Susan L. Graham, Michael A. Harrison,
and Walter L. Ruzzo**

Computer Science Division
University of California
Berkeley, California 94720

ACM Trans. Prog. Lang. Syst. 2, 3 (July 1980), 415-462.

A new algorithm for recognizing and parsing arbitrary context-free languages is presented, and several new results are given on the computational complexity of these problems. The new algorithm is of both practical and theoretical interest. It is conceptually simple and allows a variety of efficient implementations, which are worked out in detail. Two versions are given which run in faster than cubic time. Surprisingly close connections between the Cocke-Kasami-Younger and Earley algorithms are established which reveal that the two algorithms are "almost" identical. One significant use of general context-free methods is as part of a system of processing natural languages such

as English, viewing the grammar/parser as a convenient control structure for directing the analysis.

The Hearsay-II Speech-Understanding System: Integrating Knowledge to Resolve Uncertainty

Lee D. Erman

USC/Information Sciences Institute
Marina del Rey, California 90291

Frederick Hayes-Roth

The RAND Corporation
Santa Monica, California 90406

Victor R. Lesser

University of Massachusetts
Amherst, Massachusetts 01003

D. Raj Reddy

Carnegie-Mellon University
Pittsburgh, Pennsylvania 15213

Computing Surveys 12, 2 (June 1980), 213-253.

The Hearsay-II system, developed at Carnegie-Mellon University during the DARPA-sponsored five-year speech-understanding research program, represents both a specific solution to the speech-understanding problem and a general framework for coordinating independent processes to achieve cooperative problem-solving behavior. Speech-understanding, as a computational problem, reflects a large number of intrinsically interesting issues. Spoken sounds are achieved by a long chain of successive transformations, from intentions, through semantic and syntactic structuring, eventually resulting in audible acoustic waves. As a consequence, interpreting speech means effectively inverting these transformations to recover the speaker's intention from the sound. At each step in the interpretive process, ambiguity and uncertainty arise.

The Hearsay-II problem-solving framework reconstructs an intention from hypothetical interpretations formulated at various levels of abstraction. In addition, it allocates limited processing resources first to the most promising incremental actions. The final configuration of the Hearsay-II system comprises problem-solving components to generate and evaluate speech hypotheses, and a focus-of-control mechanism to identify potential actions of greatest value. Many of these specific procedures reveal novel approaches to speech problems. Most importantly, the system successfully integrates and coordinates all of these independent activities to resolve uncertainty and control combinatorics. Several adaptations of the Hearsay-II

framework have already been undertaken in other problem domains and we anticipate this trend will continue; many future systems necessarily will integrate diverse sources of knowledge to solve complex problems cooperatively.

This paper discusses the characteristics of the speech problem in particular, the special kinds of problem-solving uncertainty in that domain, the structure of the Hearsay-II system developed to cope with that uncertainty, and the relationship between Hearsay-II's structure and those of other speech-understanding systems. The paper is intended for the general computer science audience and presupposes no speech or artificial intelligence background.

Morphosemantic Analysis of -ITIS Forms in Medical Language

M.G. Pacak, L.M. Norton, and G.S. Dunham

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

Meth. Inform. Med. 19, 2 (April 1980), 99-105.

This paper describes an automated procedure for morphosemantic analysis and semantic interpretation of medical compound word forms ending in -ITIS. The requirements for morphosemantic analysis of -ITIS forms include: a) semantic classification of morphosemantic constituents forming -ITIS words forms, b) establishment of morphosemantic distribution patterns occurring in -ITIS form, c) preparation of paraphrasing rules.

Lexical Analysis of German Texts

Ingeborg Steinacker and Harald Trost

Department of Medical Cybernetics
University of Vienna

SIGLASH Newsletter 13, 3 (Sept. 1980), 6-12.

We present a computerized method for reducing inflected German words to their stem. The German language has many possibilities of inflecting words (nouns have case endings, endings of adjectives depend on case and gender of the noun they belong to, etc.) Therefore it frequently happens that an inflected word can be reduced to more than one stem. In this case all possible dictionary entries to which the word can be reduced are stored and the final selection can only be done by semantic means. The program is written in PASCAL and runs on a 16 bit machine.