

Abstracts of Current Literature

Toward a Detailed Model of Processing for Language Describing the Physical World

David L. Waltz

Coordinated Science Laboratory
University of Illinois
1101 West Springfield Avenue
Urbana, Illinois 61801

Proc. 7th IJCAI, vol. 1, August 1981, 1-6.

This paper explores the problem of judging whether or not an English sentence could correspond to a real world situation or event which is literally, physically plausible, and the related problem of representing the different possible physical situations. The judgment of plausibility can be made at a high level by checking semantic marker restrictions on verb case frame constituents. Often, however, plausibility judgement can only be based on the results of an attempt to construct (imagine) a scene that corresponds to the sentence, and which does not violate "common sense" (i.e. relevant physical laws and expected, stereotyped behavior). Methods are presented for constructing representations for different scenes which could correspond to a sentence. These methods incorporate (1) "subscripts" (sequences of scenes which comprise an event, with attached preconditions and postconditions) to express different verb senses; (2) object representations which express properties such as shape, size, weight, strength, and behavior under common conditions; (3) physical laws, encoded as constraints on behavior; (4) representation of context; and (5) robot problem solving-like methods to fit all this material together.

Language Comprehension in a Problem Solver

Douglas Wong

Department of Computer Science
Brown University
Providence, Rhode Island 02912

Proc. 7th IJCAI, vol. 1, August 1981, 7-12.

This paper describes BRUIN, a unified AI system that can perform both problem-solving and language comprehension tasks. Included in the system is a frame-based knowledge-representation language called FRAIL, a problem solving component called NASL (which is based on McDermott's problem-solving language of the same name), and a context-recognition component currently known as PRAGMATICS. The intent of this paper is to give the flavor of how the context recognizer PRAGMATICS works and what it can do. Examples are drawn from the inventory-control, restaurant and blocks-world domains.

Cancelled Due to Lack of Interest

Michael Lebowitz

Department of Computer Science
Columbia University
406 Mudd Building
New York, New York 10027

Proc. 7th IJCAI, vol. 1, August 1981, 13-15.

The parts of a typical piece of text vary greatly in interest. Presented in this paper are three ways a concept can fail to be interesting — it can be irrelevant, reconstructible, or overshadowed. The uses of interest in understanding are also discussed.

Story Generation after TALE-SPIN

Natalie Dehn

Department of Computer Science
Yale University
New Haven, Connecticut 06520

Proc. 7th IJCAI, vol. 1, August 1981, 16-18.

TALE-SPIN, the last major AI attempt at story generation, approached the problem of making up stories primarily from the perspective of an impartial world simulator. *AUTHOR* is a program (under development) which generates stories as a creative reasoner in pursuit of her own narrative goals. It is thus intended to simulate an author's mind as she makes up a story, rather than the world as things happen in it. The four major forces driving the story generation process, according to the *AUTHOR* model, are author intentionality, conceptual reformation, reminding, and the opportunity enhancement metagoal.

Modeling Informal Debates

Rachel Reichman

Elect. Engrg. and Comp. Sci., C-014
University of California
La Jolla, California 92093

Proc. 7th IJCAI, vol. 1, August 1981, 19-24.

Many rules of formal debate are well documented, are of common knowledge, and are "looked-up" in preparation for planned debating. Informal debates, on the other hand, are highly dynamic, are complex, and are spontaneously generated with no prior rule-book preparation. They too, however, are rule-governed. In this paper I present an abstract process model capable of modeling "well-formed" argument structures that occur in ordinary conversations. The formalization rests on a general theoretical framework for discourse engagement encapsulated in a discourse ATN grammar. A major feature of the system is its segmentation of discourse utterances into functionally related context spaces.

A Knowledge-Based Approach to Language Processing: A Progress Report

Robert Wilensky

Department of EECS
University of California
Berkeley, California 94720

Proc. 7th IJCAI, vol. 1, August 1981, 25-30.

We present a model of natural language use meant to encompass the language-specific aspects of understanding and production. The model is motivated by the pervasiveness of non-generative language, by the desirability of a language analyzer and a language production mechanism to share their knowledge, and by the advantages of knowledge engineering features such as ease of extension and modification.

This model has been used as the basis for PHRAN, a language analyzer, and PHRED, a language production mechanism. We have implemented both these systems using a common knowledge base; we have produced versions of PHRAN that understand Spanish and Chinese by only changing the knowledge base and not modifying the program; and we have implemented PHRAN using the query language of a conventional relational data base system, and compared the performance of this system to a conventional LISP implementation.

The Need for Referent Identification as a Planned Action

Philip R. Cohen

Department of Computer Science
Oregon State University
Corvallis, Oregon 97331

Proc. 7th IJCAI, vol. 1, August 1981, 31-36.

The paper presents evidence that speakers often attempt to get hearers to identify referents as a separate step in the speaker's plan. Many of the communicative acts performed in service of such referent identification steps can be analyzed by extending a plan-based theory of communication for task-oriented dialogues to include an *action* representing a hearer's identifying the referent of a description — an action that is reasoned about in speakers' and hearers' plans. The phenomenon of addressing referent identification as a separate goal is shown to distinguish telephone from teletype task-oriented dialogues and thus has implications for the design of speech-understanding systems.

Integration, Unification, Reconstruction, Modification: An Eternal Parsing Braid

Michael G. Dyer

Department of Computer Science
Yale University
New Haven, Connecticut 06520

Proc. 7th IJCAI, vol. 1, August 1981, 37-42.

BORIS is an integrated natural language understanding system for narratives. In an integrated system, processes of event assimilation, inference, and episodic memory search occur on a word-by-word basis as parsing proceeds. "Parsing" here refers to the task of building a conceptual representation for each natural language expression. In addition to being integrated, the BORIS parser is also a unified parser. The same parser is used both at story understanding time and question answering time. This paper explores some of the consequences which arise when the same parser serves both tasks. For instance, one such consequence is that BORIS often knows the answer to a question before it has completely understood the question.

Design Characteristics of a Machine Translation System

M. King

ISSCO
Université de Genève
17 Rue de Candolle
CH-1205 Genève, SWITZERLAND

Proc. 7th IJCAI, vol. 1, August 1981, 43-46.

This paper distinguishes a set of criteria to be met by a machine translation system (EUROTRA) currently being planned under the sponsorship of the Commission of the European Communities and attempts to show the effect of meeting those criteria on the overall system design.

High Level Memory Structures and Text Coherence in Translation

C.J. Yang

Department of Computer Science
Yale University
New Haven, Connecticut 06520

Proc. 7th IJCAI, vol. 1, August 1981, 47-49.

Various memory organization schemes have been proposed in the last five years. Lots of intelligent computer systems have been experimenting with memory schemes like scripts, plans, goals, and MOPs in the domain of text understanding and information retrieval. In this paper, the focus is on the problem of translating sentences that involve lexical items which do not have equivalent counterparts in the target language. Examples are drawn from translations between English and Mandarin Chinese.

Natural Lang. Dialogue about Moving Objects in an Automatically Analyzed Traffic Scene

H. Marburger, B. Neumann, and H-J. Novak

Fachbereich Informatik
Universitaet Hamburg
Schlueterstrasse 70
D-2000 Hamburg 13

Proc. 7th IJCAI, vol. 1, August 1981, 49-51.

This contribution is concerned with natural language dialogues about scenes with moving objects. Two systems are connected: a natural language dialogue system originally conceived for static scenes and an emerging scene analysis system for real-world TV-frame sequences. The latter produces time dependent object descriptions which serve as a referential database for inquiries. The time intervals relevant for answering the questions are determined from domain specific parameters, the context of the dialogue, the tense of the verbs and time adverbials. For checking the correspondence between a verbally specified motion and a trajectory, predicates are evaluated which can be deduced from the verb's case-frame.

Using Language and Context in the Analysis of Text

Yigal Arens

Department of EECS
University of California, Berkeley
Berkeley, California 94720

Proc. 7th IJCAI, vol. 1, August 1981, 52-57.

We describe a theory of natural language understanding within which we identify two separate components, a *language centered* one and a *context centered* one. The former component uses a knowledge base consisting of pairings of phrases with the concepts associated with them to determine the meaning of utterances. The latter component clarifies the meaning found by the first one and makes it more specific by attempting to reconcile it with the context of the utterance. We have constructed a program called PHRAN (PHRasal ANalyzer) which performs the task of the language centered component.

Opportunistic Processing in Arguments

Rod McGuire, Lawrence Birnbaum, Margot Flowers

Department of Computer Science
Yale University
New Haven Connecticut 06520

Proc. 7th IJCAI, vol. 1, August 1981, 58-60.

In two previous papers we have proposed a part of a computational theory of argumentation, including representations for argument structure and rules for using those representations in understanding and in rebutting. One property of the model which we emphasized is the way in which argument mechanisms and inferential memory can each help to direct the

processing of the other. In particular, we presented examples in which inferential memory can uncover good rebuttals to an input as a side-effect of the processing that naturally goes on in trying to understand that input. When such opportunities for rebuttal are noticed during understanding, they render unnecessary the use of argument rules to find a response, since one has already been discovered.

Natural Language Interaction with Dynamic Knowledge Bases: Monitoring as Response

Eric Mays, Sitaram Lanka, Aravind K. Joshi, and Bonnie L. Webber

Department of Computer and Information Science
University of Pennsylvania
Philadelphia, Pennsylvania 19104

Proc. 7th IJCAI, vol. 1, August 1981, 61-63.

In this communication, we discuss an interesting aspect of natural language interaction with dynamically changing knowledge bases — the ability to monitor for relevant future changes in that knowledge. We also indicate the status of our current work in this area and the overall goals of our research on question-answering and monitoring dynamic knowledge bases.

Variable-depth Natural Language Understanding

Daniel Kayser

E.R.A. 452 du CNRS
Laboratoire de Recherche en Informatique
Bât.490 - Campus d'Orsay
91405 Orsay, FRANCE

Daniel Coulon

L.A. 262 du CNRS
Centre de Recherche en Informatique de Nancy
ENSMIM - Parc de Saurupt
54042 Nancy, FRANCE

Proc. 7th IJCAI, vol. 1, August 1981, 64-66.

Standard AI representations of knowledge operate at fixed depth (i.e., the objects manipulated are described by an amount of information which remains constant for every task). Contrary to this approach, Variable Depth Processing (VDP) uses a *progressive description* of objects, tries *different strategies* according to the quality of the result it needs, and continually controls this quality by means of an *evaluation of the approximations* it makes. Contextual Production Rules are shown to be an effective way to implement some features of VDP. We are currently developing a VDP question-answering system which works on texts concerning a non-technical subject, namely an excerpt of a general public-oriented encyclopaedia.

Generalizations Based on Explanations

Gerald DeJong

Coordinated Science Laboratory
University of Illinois
1101 West Springfield Avenue
Urbana, Illinois 61801

Proc. 7th IJCAI, vol. 1, August 1981, 67-70.

This paper describes a new project in computer learning. The phenomenon under study is a kind of "insight learning" of procedural schemata. The system described here is designed to grasp some principle underlying a natural language input. The underlying principle results in a new schema for the system. Once acquired, the schema serves the same purpose as the other schemata in the system: it aids in processing future natural language inputs.

The process that the system uses is called *explanatory schema acquisition*. The basic idea behind it is that the causal connections in an understood representation of a new input can be used to propose and propagate constraints on slot fillers. That is, from one particular instance or situation the system can "reason out" the general structure underlying that instance. The system is therefore capable of learning from just one example.

Viewing Word Expert Parsing as Linguistic Theory

Steven Small

Department of Computer Science
Mathematical Sciences Building
University of Rochester
Rochester, New York 14627

Proc. 7th IJCAI, vol. 1, August 1981, 70-76.

The *Word Expert Parser* is a computer program that analyzes fragments of natural language text in order to extract their meaning in context. The construction of the program has led to the development of a linguistic theory based on notions orthogonal to those traditionally found at the heart of such theories. Word Expert Parsing explains the understanding of textual fragments containing highly idiosyncratic elements, such as idioms, collocations, clichés, and colligations, as well as lexical sequences that contain interesting structural phenomena. The theory perceives the individual word of language as the organizing unit for linguistic knowledge, and views understanding as consisting of *lexical interactions* among procedural word experts. This paper describes four classes of lexical interaction required to explain the understanding of sentences in context, *idiosyncratic interaction*, *linguistic interaction*, *discourse interaction*, and *logical interaction*. The paper purposely avoids programming details in order to focus on Word Expert Parsing as linguistic theory.

A Plot Understanding System on Reference to Both Image and Language

Norihiro Abe, Itsuya Soga, and Saburo Tsuji

Department of Control Engineering
Osaka University
Toyonaka, Osaka, JAPAN

Proc. 7th IJCAI, vol. 1, August 1981, 77-84.

A system is described that can understand a plot of a story on reference to both image and linguistic information. As input, a series of line drawings with colors and narrations in English concerning these drawings are given to the system. It searches the objects suggested to be in the scene by the narrations, finding relations among them, making the world model by using its world knowledge. Its reference to those drawings makes it easy for the system to analyze complicated structures in the narration sentences, such as those of prepositions, and guides the process reasoning about the CD representation using rules and demons. At the end of this paper, a result on QA is shown.

Metaphor Interpretation as Selective Inferencing

Jerry R. Hobbs

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

Proc. 7th IJCAI, vol. 1, August 1981, 85-91.

Metaphor pervades natural language discourse. This paper describes a computational approach to the interpretation of metaphors. It is based on a natural language processing system that uses the discourse problems posed by a text to select the relevant inferences. The problem of interpreting metaphors can then be translated into the problem of selecting the relevant inferences to draw from the metaphorical expression. Thus, a metaphor is frequently given a correct interpretation as a by-product of the other things a natural language system has to do. Two examples of metaphors are given — a spatial metaphor schema from computer science, and a novel metaphor — and it is shown how the interpretation problem for each can be translated into a selective inferencing problem and solved by the ordinary operations of the system. This framework sheds light on the analogical processes that underlie metaphors and begins to explain the power of metaphor.

A Computer Model of Child Language Acquisition

Mallory Selfridge

Department of EE and CS
University of Connecticut
Storrs, Connecticut 06268

Proc. 7th IJCAI, vol. 1, August 1981, 92-96.

Children learn different aspects of language in a characteristic order, and make characteristic errors during acquisition. This paper focuses on five specific data, and explains these data in terms of two hypotheses regarding the relationship of comprehension to generation and the relationship between meaning and syntax. A computer model, CHILD, is described which embodies these hypotheses and which manifests the same five data. CHILD'S performance suggests that the explanations are plausible.

A Theory of Language Acquisition Based on General Learning Principles

John R. Anderson

Department of Psychology
Carnegie-Mellon University
Pittsburgh, Pennsylvania 15213

Proc. 7th IJCAI, vol. 1, August 1981, 97-103.

A simulation model is described for the acquisition of the control of syntax in language generation. This model makes use of general learning principles and general principles of cognition. Language generation is modeled as a problem-solving process involving principally the decomposition of a to-be-communicated semantic structure into a hierarchy of subunits for generation. The syntax of the language controls this decomposition. It is shown how a sentence and semantic structure can be compared to infer the decomposition that led to the sentence. The learning processes involve generalizing rules to classes of words, learning by discrimination the various contextual constraints on a rule application, and a strength process which monitors a rule's history of success and failure. This system is shown to apply to the learning of noun declensions in Latin, relative clause constructions in French, and verb auxiliary structures in English.

Inductive Learning of Pronunciation Rules by Hypothesis Testing and Correction

S. Oakey and R.C. Cawthorn

Department of Computer Science
Teesside Polytechnic
Middlesbrough, Cleveland, ENGLAND

Proc. 7th IJCAI, vol. 1, August 1981, 109-114.

This paper describes a system that learns the rules of pronunciation inductively. It begins with a set of 26 rules for single-letter pronunciation. Individual words are presented to it, and the system uses its rule set to hypothesize a pronunciation. This is compared

with a dictionary pronunciation, and if any part of the pronunciation is incorrect new rules are created to handle the word as an exception condition.

These rules are checked for similarity with others already produced, and where suitable a "general" rule is produced to deal with two or more created rules. The effect is to produce rules that are more and more general, and these approach the general pronunciation rule sets that have been produced manually by other workers.

Summarizing Narratives

Wendy G. Lehnert, John B. Black, and Brian J. Reiser

Department of Computer Science
Yale University
New Haven, Connecticut 06520

Proc. 7th IJCAI, vol. 1, August 1981, 184-189.

Most research on narrative text summarization has been conducted within the paradigm of experimental psychology. But recent language processing research in artificial intelligence suggests that the predominant theory of text summarization requires further examination. Seemingly minor structural modifications of a story can result in significant alterations of summary behavior. In this paper, highlights of summary data from 72 subject are presented and analyzed in terms of two competing summarization models: (1) the story grammar model of psychology, and (2) the plot unit model developed in artificial intelligence. We show how selected story grammar predictions compare to the plot unit predictions for short term summarization and then identify two complicating factors that have a major impact on summarization behavior.

Text Plans and World Plans in Natural Discourse

Jerry R. Hobbs

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

Michael Agar

Department of Anthropology
University of Maryland
College Park, Maryland 20742

Proc. 7th IJCAI, vol. 1, August 1981, 190-196.

Discourse is both *about* the world and an accomplishment *in* the world. This fact has led to two approaches to the study of discourse in artificial intelligence: one investigating "text plans," the other, "world plans." By analyzing a fragment of a narrative in which both kinds of plans figure importantly, we explore the relationship between the two kinds of plans, looking toward a synthesis of the two approaches.

Recognizing Intended Meaning and Speakers' Plans

Candace L. Sidner and David J. Israel

Bolt Beranek and Newman Inc.

10 Moulton Street

Cambridge, Massachusetts 02238

Proc. 7th IJCAI, vol. 1, August 1981, 203-208.

Human conversational participants depend upon the ability of their partners to recognize their intentions, so that those partners may respond appropriately. In such interactions, the speaker can encode his intentions that the hearer act in a variety of sentence types. Instead of telling the hearer what to do, the speaker may just state his goals, and expect a response that meets these goals. This paper presents a new model for recognizing the speaker's intended meaning in determining a response. We show that this recognition makes use of the speaker's plan, his beliefs about the domain and about the hearer's relevant capacities.

Character Tracking and the Understanding of Narratives

Brian J. Reiser

Cognitive Science Program

Yale University

New Haven, Connecticut 06520

Proc. 7th IJCAI, vol. 1, August 1981, 209-212.

Recent work on the understanding of natural language narratives has emphasized representations composed of goals, plans, and their outcomes. A problem that has received little attention however is the influence of *perspective* in understanding a narrative. Perspective may operate on many levels. Understanding proceeds as a focused tracking of the fate of a character's goals. When attention is focused on a character in a narrative, each new input is then evaluated from that character's perspective. Processing the story from one character's perspective is one way that the inferencing process might be constrained. Less relevant inferences need never be made or integrated into the representation.

Representing Implicit and Explicit Time Relations in Narrative

Lynette Hirschman and Guy Story

Linguistic String Project

New York University

251 Mercer Street

New York, New York 10012.

Proc. 7th IJCAI, vol. 1, August 1981, 289-295.

This paper describes a representation for time relations in narrative. The time relations are based on both explicit sources of time information (e.g., adverbial expressions or tense) and implicit sources, such as multiple reference to a single event, narrative time progression and earlier events implied by change of

state words. Natural language processing is used to analyze the input text into a set of subject-verb-object units, connected by binary connectives; these units correspond to the events of the narrative. With each event is associated a time in relation to another event, adjusted by an optional time quantity. These time relations have a natural representation as a directed graph whose nodes are time points and whose edges are time intervals. The algorithm for extracting the time relations from a text is illustrated for an excerpt from a hospital discharge summary.

The Nature of Generalization in Understanding

Michael Lebowitz

Department of Computer Science

Columbia University

406 Mudd Building

New York, New York 10027

Proc. 7th IJCAI, vol. 1, August 1981, 348-353.

True understanding of natural language text requires the inclusion of generalization and long-term memory. This paper describes the generalization process and memory used in the Integrated Partial Parser (IPP), a computer program that reads and remembers news stories. The need for generalization and generalization-based memory as an integral part of understanding natural language text is illustrated with examples from IPP. In addition, the nature of generalization is discussed.

Directing and Re-directing Inference Pursuit: Extra-textual Influences on Text Interpretation

Richard H. Granger Jr.

Department of Information and Computer Science

University of California

Irvine, California 92717

Proc. 7th IJCAI, vol. 1, August 1981, 354-361.

Understanding a text depends on a reader's ability to construct a coherent interpretation that accounts for the statements in the text. However, a given text does not always imply a unique coherent interpretation. In particular, readers can be steered away from an otherwise plausible explanation for a story by such extra-textual factors as the source of the text, the reading purpose, interruptions during reading, or repeated re-questioning of the reader. Some of these effects have been observed in experiments in cognitive psychology. This paper presents a computer program called MACARTHUR that can vary both the depth and direction of its inference pursuit in response to re-questioning, resulting in a series of markedly different interpretations of the same text.

Towards Automating Explanations

R.E. Cullingford, M.W. Krueger, M. Selfridge, and
M.A. Bienkowski

Department of EE and CS
University of Connecticut
Storrs, Connecticut 06268

Proc. 7th IJCAI, vol. 1, August 1981, 362-367.

This paper discusses an approach to the modelling of the explanation process within the framework of a graphics-based CAD system currently under development, which can describe its own use, including the common ways to make and recover from errors. With a coordinated textual and pictorial display, the system, CADHELP, simulates an expert *demonstrating* the operation of the graphical features of the CAD tool. It consults a knowledge base of *feature scripts*, built up using situational script and commonsense algorithmic methods, to explain a feature, generate prompts as the feature is being operated, and to give certain types of "help" when a feature is misused. CADHELP provides these services by summarizing the feature script in different ways depending upon what it has told the user previously. The summarization process is based upon a series of "sketchification" strategies, which prescribe which parts of a knowledge structure, a causal chain, or a single concept can be thrown away, since the listener should be able to infer them.

Using Active Connection Graphs for Reasoning with Recursive Rules

Donald P. McKay and Stuart C. Shapiro

Department of Computer Science
State University of New York at Buffalo
Amherst, New York 14226

Proc. 7th IJCAI, vol. 1, August 1981, 368-374.

Recursive rules, such as "Your parents' ancestors are your ancestors," although very useful for theorem proving, natural language understanding, question-answering and information retrieval systems, present problems for many such systems, either causing infinite loops or requiring that arbitrarily many copies of them be made. SNIP, the SNePS Inference Package, can use recursive rules without either of these problems. A recursive rule causes a cycle to be built in an *active connection graph*. Each pass of data through the cycle results in another answer. Cycling stops as soon as either the desired answer is produced, no more answers can be produced, or resource bounds are exceeded.

Control of Inference: Role of Some Aspects of Discourse Structure - Centering

Aravind K. Joshi

Department of Computer and Information Science
University of Pennsylvania
Philadelphia, Pennsylvania 19104

Scott Weinstein

Department of Philosophy
University of Pennsylvania
Philadelphia, Pennsylvania 19104

Proc. 7th IJCAI, vol. 1, August 1981, 385-387.

The purpose of this communication is to examine one particular aspect of discourse structure, namely, a discourse construct called *center* of a sentence (utterance) *in discourse* and its relation to the larger issue of control of inference. We describe very briefly the notion of center(s) of a sentence in discourse and discuss how the centering phenomenon might be incorporated in a formal model of inference and its relation to the intrinsic complexity of certain inferences.

The Design and an Example Use of Hearsay-III

Lee Erman, Philip London, and Stephen Fickas

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

Proc. 7th IJCAI, vol. 1, August 1981, 409-415.

Hearsay-III provides a framework for constructing knowledge-based expert systems. While Hearsay-III makes no commitment to any particular application domain, it does supply a variety of generally applicable facilities. These include representation primitives and an interpreter for large-grained, flexibly schedulable production rules called knowledge sources. A detailed overview of the motivations behind Hearsay-III and the facilities it provides are presented. Finally, an application of Hearsay-III is described.

Anaphora for Limited Domain Systems

Philip J. Hayes

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

Proc. 7th IJCAI, vol. 1, August 1981, 416-422.

This paper presents a simple mechanism for the resolution of anaphora in limited domain natural language systems. For such domains, this mechanism provides functionality equivalent to the natural communication mechanism of anaphora as used and understood by people, but without the deep inferencing or cognitive modeling required for full simulation of human performance. The mechanism covers simple pronoun anaphora, and set selection anaphora (e.g. "last one," "one before," "others"). It was developed to provide the most efficient and effective communication between system and user, even if this meant di-

verging significantly from human performance when this performance was impractical to reproduce. In cases of radical divergence, we were careful to make the behaviour of the mechanism very simple and easy to predict. In this way, the user can rely either on his experience of human performance or his knowledge of the artificial, but simple, substitute to predict the behaviour of the system in response to his inputs. An algorithmic description of an implemented version of the mechanism is presented. A similar approach to other aspects of man-machine interfaces is recommended as a promising way to address the problem of habitability that still plagues all natural language computer interfaces.

Figuring out What the User Wants - Steps toward an Automatic Yellow Pages Assistant

Anatole Gershman

Schlumberger-Doll Research
P.O. Box 307
Ridgefield, Connecticut 06877

Proc. 7th IJCAI, vol. 1, August 1981, 423-425.

An experimental system, AYP, for automatic Yellow Pages assistance, is described. The system, which operates in the domain of automobiles, automobile parts, and related objects, reads the user's request in simple English, analyzes it and represents it in terms of the system's conceptual primitives. From this, the system tries to figure out the intent of the request and formulate a Yellow Pages query. It paraphrases the request back to the user in English and searches its data base for the relevant Yellow Pages categories. The system serves as a research vehicle for experiments with its various components and user interfaces.

Computing a Map from Michi-Annai-Bun or Written Directions

Teiji Furugori

University of Electro-communications
Chofu, Tokyo, JAPAN

Proc. 7th IJCAI, vol. 1, August 1981, 426-428.

This paper describes processes of transforming michi-annai-bun, literally street-guide-sentence, into a map on a display device: we first show how such transformation is performed, and then we discuss what it means in terms of understanding natural language expressions.

GLP: A General Linguistic Processor

G. Goerz

University of Erlangen-Nuernberg
RRZE Martensstr.1
D-8520 Erlangen, WEST GERMANY

Proc. 7th IJCAI, vol. 1, August 1981, 429-431.

GLP is a general linguistic processor for the analysis and generation of natural language, based on a second generation version of the General Syntactic Processor of Kaplan and Kay. It is part of a speech understanding system currently under development at the Computer Science Department of our university.

Multi-Strategy Construction-Specific Parsing for Flexible Data Base Query and Update

Philip J. Hayes and Jaime G. Carbonell

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

Proc. 7th IJCAI, vol. 1, August 1981, 432-439.

The advantages of a multi-strategy, construction-specific approach to parsing in applied natural language processing are explained through an examination of two pilot parsers we have constructed. Our approach exploits domain semantics and prior knowledge of expected constructions, using multiple parsing strategies each optimized to recognize different construction types. It is shown that a multi-strategy approach leads to robust, flexible, and efficient parsing of both grammatical and ungrammatical input in limited-domain, task-oriented, natural language interfaces. We also describe plans to construct a single, practical, multi-strategy parsing system that combines the best aspects of the two simpler parsers already implemented into a more complex, embedded-constituent control structure. Finally, we discuss some issues in data base access and update, and show that a construction-specific approach, coupled with a case-structured data base description, offers a promising approach to a unified, interactive data base query and update system.

A Deterministic Analyzer for the Interpretation of Natural Language Commands

Leonardo Lesmo, Daniela Magnani, and Piero Torasso

Instituto di Scienze dell'Informazione
Università di Torino
C.so Massimo D'Azeglio
42 - 10125 Torino, ITALY

Proc. 7th IJCAI, vol. 1, August 1981, 440-442.

This paper describes a system which translates a query in the Italian language into a representation which can be immediately interpreted as a sequence of algebraic operations on a relational data base. The use of a lookahead buffer allows the system to operate deterministically. Different knowledge sources are used to cope with semantics (associated with the lexicon) and syntax (represented as pattern-action rules). These knowledge sources cooperate during the query translation so that independent translations of the command are avoided. Therefore, the term "determinism" is used to mean that all the structures

built during the process concur to build the final command representation.

A General Semantic Analyzer for Data Base Access

B.K. Boguraev and K. Sparck Jones

Computer Laboratory
University of Cambridge
Corn Exchange Street
Cambridge CB2 3QG, ENGLAND

Proc. 7th IJCAI, vol. 1, August 1981, 443-445.

The paper discusses the design principles and current status of a natural language front end for access to data bases. This is based on the use, first, of a semantically-oriented question analyzer exploiting general, language-wide semantic categories and patterns, rather than data base-specific ones: and, second, of a data base-oriented translation component for obtaining search specifications from the meaning representations for questions derived by the analyzer. This approach is motivated by the desire to reduce the effort of providing data base-specific material for the front end, by the belief that a general analyzer is well suited to the "casual" data base user, and by the assumption that the rich semantic apparatus used will be both adequate as a means of analysis and appropriate as a tool for linking the characterizations of input and data language items. The paper describes this approach in more detail, with emphasis on the existing, tested, analyzer.

A Metalanguage Representation of Relational Databases for Deductive Q-A Systems

Kurt Konolige

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

Proc. 7th IJCAI, vol. 1, August 1981, 496-503.

This paper presents a method of formally representing the information that exists in a relational database. The primary utility of such a representation is for deductive question-answering systems that must access an existing relational database. To respond intelligently to user inquiries, such systems must have a more complete representation of the domain of discourse than is generally available in the database. The problem that then arises is how to reconcile the information present in the database with the domain representation so that database queries can be derived to answer the user's inquiries. Here we take the formal approach of describing a relational database as the model of a first-order language. Another first-order language, the metalanguage, is used both to represent the domain of discourse, and to describe the relationship of the database to the domain. This view proves

particularly useful in two respects. First, by axiomatizing the database language and its associated model in a metatheory, we are able to describe in a powerful and flexible manner how the database corresponds to the domain of discourse. Secondly, viewing the database as a mechanizable model of the database language enables us to take advantage of the computational properties of database query language processors. Once a database query that is equivalent to an original query is derived, it can be evaluated against the database to determine the truth of the original query. Thus the algebraic operations of the database processor can be incorporated in an elegant way into the deductive process of question-answering.

Explaining and Justifying Expert Consulting Programs

William R. Swartout

Laboratory for Computer Science
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139

Proc. 7th IJCAI, vol. 2, August 1981, 815-823.

Traditional methods for explaining programs provide explanations by converting to English the code of the program or traces of the execution of that code. While such methods can provide adequate explanations of what the program does or did, they typically cannot provide justifications of the code without resorting to canned-text explanations. That is, such systems cannot tell why what the system is doing is a reasonable thing to be doing. The problem is that the knowledge required to provide these justifications is needed only when the program is being written and does not appear in the code itself.

The XPLAIN system uses an automatic programmer to generate the consulting program by refinement from abstract goals. The automatic programmer uses a domain model, consisting of facts about the application domain, and a set of domain principles which drive the refinement process forward. By examining the refinement structure created by the automatic programmer it is possible to provide justifications of the code. This paper discusses the system described above and outlines additional advantages this approach has for explanation.

Last Steps towards an Ultimate PROLOG

A. Colmerauer, H. Kanoui, and M. Van Caneghem

Groupe d'Intelligence Artificielle
Faculte des Sciences de Luminy
Universite d'Aix-Marseille II
13288 Marseille Cedex 9, FRANCE

Proc. 7th IJCAI, vol. 2, August 1981, 947-948.

A portable version of PROLOG, an Artificial Intelligence language, is presented. A complete system has been implemented on a micro-computer using a floppy-disk virtual memory. The general methodology of the implementation is discussed in terms of an abstract machine (Micromegas) supporting a language (Candide) in which the PROLOG system is written.

Six Topics in Search of a Parser: An Overview of AI Language Research

Eugene Charniak

Department of Computer Science
Brown University
Providence, Rhode Island 02912

Proc. 7th IJCAI, vol. 2, August 1981, 1079-1087.

My purpose in this paper is to give an overview of natural language understanding work within artificial intelligence (AI). I will concentrate on the problem of parsing — going from natural language input to a semantic representation. Naturally, the form of semantic representation is a factor in such discussions, so it will receive some attention as well. Furthermore, I doubt that parsing can be completely isolated from text processing issues, and hence I will touch upon such seemingly non-parsing issues as script application. Nevertheless, the topic is parsing.

Unfortunately, to present AI parsing work with any sort of historical accuracy would be to produce a bewildering forest of names (both of people and programs). Rather, I will try to extract from the historical record a group of ideas which I believe can be molded into a coherent framework. Naturally, even within these limitations my remarks will be sketchy — this is an article, not a book.

Center-Embedding Revisited

Michael B. Kac

Department of Linguistics
University of Minnesota
Minneapolis, Minnesota 55455

Proc. 3rd Ann. Conf. Cog. Sci. Soc., Aug. 1981, 123-125.

The severe comprehension difficulty associated with certain center-embedding constructions is perhaps the best known of psychosyntactic phenomena. Most attempts at explanation have been variations on a single theme — that the c.e. configuration leads to an overload of short-term memory during processing. That this is not the whole story can be seen from considering the fact, rarely noted, that c.e. constructions exist which are understood quite easily. Several examples are discussed.

The Natural Natural Language Understander

Henry Hamburger

Computer Science Section
National Science Foundation
1800 G Street, N.W.
Washington, D.C. 20550

Stephen Crain

Department of Computer Sciences
The University of Texas
Austin, Texas 78712

Proc. 3rd Ann. Conf. Cog. Sci. Soc., Aug. 1981, 128-130.

This study of natural language comprehension by natural understanding systems (children) is based on a procedural analysis represented in the form of a programming language. To clarify what is cognitively required for a child to respond appropriately to certain expressions in English, we show how these forms can be translated into procedures in a high-level programming language. It is then possible to discuss two kinds of difficulties a natural language form can present to the listener: (1) incompatibility of the form with its associated procedure, and (2) complexity of that procedure. An example of procedure complexity is the nesting of loops, whereas a contributor to incompatibility is a word or contiguous phrase that corresponds to separated pieces of the procedure. We present evidence of both types of difficulty from experiments with children and compare the predictions of our procedural view with those of a less detailed syntactic explanation that has been advanced for a subset of the phenomena.

Why Do Children Say "Goed"? A Computer Model of Child Language Generation

Mallory Selfridge

Department of EE and CS
University of Connecticut
Storrs, Connecticut 06268

Proc. 3rd Ann. Conf. Cog. Sci. Soc., Aug. 1981, 131-133.

An important question in modelling child language generation is why children say regular forms of irregular words, such as "goed," during development, although they never hear them. Three other general characteristics of children's generation also require explanation. First, Benedict's work suggests clearly that comprehension of various aspects of language precede the generation of those aspects. Second, the length of the utterances children say become generally longer as development proceeds. Third, Wetstone and Friedlander suggest that first children say things in the wrong order, and then say things in the correct order.

In order to address these issues, this paper explores the hypothesis that learning to talk is driven by learning to understand. This hypothesis begins by assuming that the principal effect of learning to understand is

the development of the lexicon as additional words are learned and their "definitions" are refined and modified. It further assumes that the language generation process is not learned, but is an innate part of a child's cognitive repertoire. Finally, it states that the ability to generate grows as the lexicon develops during the development of comprehension. The hypothesis predicts that a computer model which incorporated it would display the characteristics described above.

Writing with a Computer

Ira Goldstein

Xerox Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, California 94304

Proc. 3rd Ann. Conf. Cog. Sci. Soc., Aug. 1981, 145-148.

This essay conjectures that an author's planning process will be facilitated by a tool that represents his plan at various levels of abstraction as a network of subgoals, with the subgoal not necessarily restricted to a linear order. Machine reasoning on such structures has been explored in artificial intelligence research. Our proposal is to make these structures available to the writer as a calculus for representing his essays and to use the computer as an interactive editing tool to manipulate them.

MOPs and Learning

Roger C. Schank

Department of Computer Science
Yale University
New Haven, Connecticut 06520

Proc. 3rd Ann. Conf. Cog. Sci. Soc., Aug. 1981, 166-170.

This paper is an attempt to sketch out some of what MOPs are about. A MOP is an orderer of scenes. A scene is a memory structure that groups together actions with a common goal, a common time, and some other common thread. It provides a sequence of very general actions. Specific memories are stored in scenes, indexed with respect to how they differ from the general action in the scene. Scenes actually point to specific memories. MOPs do not. MOPs merely point to scenes. Scripts are particularly common instantiations of scenes.

Shaping Explanations: Effects of Questioning on Text Interpretation

Richard H. Granger Jr.

Computer Science Department
University of California
Irvine, California 92717

Proc. 3rd Ann. Conf. Cog. Sci. Soc., Aug. 1981, 193-197.

Results in cognitive psychology have shown that readers can be steered away from an otherwise plausible interpretation of a story by extra-textual factors

such as the source of the text, the stated reading purpose, interruptions and repetition of questions about the text. For instance, successive repetitions of the same question about a given text will often elicit a series of alternative interpretations of the text. This effect cannot be accounted for by established principles of text processing behavior, such as people's preference for cohesive and parsimonious representations of text. This paper presents a computer program called MACARTHUR, which models this behavior by varying the depth and direction of its inference pursuit in response to re-questioning, resulting in a series of markedly different interpretations of the same text. In light of the results, some new experiments are suggested in hopes of arriving at a new principle, beyond cohesion and parsimony, to account for the observed text processing behavior.

Memory in Story Invention

Natalie Dehn

Department of Computer Science
Yale University
New Haven, Connecticut 06520

Proc. 3rd Ann. Conf. Cog. Sci. Soc., Aug. 1981, 213-216.

AUTHOR is a story generating program (under development) being built as a model of how human authors make up stories. Like TALE-SPIN, AUTHOR requires human-like knowledge of the world, but unlike TALE-SPIN, AUTHOR also requires human-like memory organization of this knowledge. The two features of human memory most essential to the AUTHOR model of story generation are: (1) reconstruction, and (2) reminding. The former is responsible for the directed nature of making up stories, the latter for the author's more "fortuitous" ideas and insights.

Controlling Parsing by Passing Messages

Brian Phillips and James Hendler

Texas Instruments
P.O. Box 225936, MS 371
Dallas, Texas 75265

Proc. 3rd Ann. Conf. Cog. Sci. Soc., Aug. 1981, 228-231.

The functional segmentation of linguistic knowledge into rules about form and rules about meaning has been vital in unravelling the complexities of language. However, it does not follow that the process of analysis will respect the same boundaries, and so the very segmentation that provided the insights can be troublesome when one seeks to create a dynamic model of language. We believe that a language understanding system should have the ability to bring syntactic and semantic knowledge to bear on the analysis at many points in the computation. This enables it to resolve the alternatives as soon as possible and prevent the flow of extraneous analyses to later phases. Our approach to creating such a model is to use the notion of

a society of communicating, knowledge-based problem-solving experts, called "actors". These actors can communicate by passing messages to any other actor in the system. This flexible control structure allows actors at any level of the analysis to interact with actors at other levels.

A Parser with Something for Everyone

Eugene Charniak

Department of Computer Science
Brown University
Providence, Rhode Island 02912

Proc. 3rd Ann. Conf. Cog. Sci. Soc., Aug. 1981, 231-234.

We present a syntactic parser, Paragram, which tries to accommodate three goals. First, it will parse, in a natural way, ungrammatical sentences. Secondly, it aspires to "capture the relevant generalizations," as in transformational grammar, and thus its rules are in virtual one-to-one correspondence with typical transformational rules. Finally, it promises to be reasonably efficient, especially given certain limited parallel processing capabilities.

GLISP: An Efficient, English-Like Programming Language

Gordon S. Novak Jr.

Department of Computer Science
University of Texas
Austin, Texas 78712

Proc. 3rd Ann. Conf. Cog. Sci. Soc., Aug. 1981, 249-252.

My earlier research on computer understanding of physics problems, stated in English, has convinced me that English is best viewed as a programming language. That is, an English sentence does not contain the message to be transmitted to the reader, but rather is a program which provides the minimum information necessary for the reader to construct the message from what the reader already knows. Study of the ways in which English permits compact expression of complex ideas reveals several features which would be useful if incorporated into programming languages. GLISP is a LISP-based programming language which permits English-like programs containing definite references.

Computerized Language Processing for Multiple Use of Narrative Discharge Summaries

**Naomi Sager, Lynette Hirschman,
and Margaret Lyman**

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

Proc. 2nd Ann. Symp. on Comp. Appl. in Med. Care, IEEE, 1978, 330-343.

At New York University, computer programs have been developed that convert natural language medical records into a structured data base, i.e., into a table containing the same information as the stored documents. In this form specific information can be quickly retrieved, and summaries of the different kinds of information in the documents can be automatically generated. The automatic conversion of the information from its free-text form to a tabular form is called information formatting. This paper describes the application of the information formatting programs to a small set of pediatric discharge summaries for hospitalizations due to sickle cell disease. The programs created a table of approximately 50 columns in which each different type of information in the documents appeared under a separate heading. From this, a retrieval program extracted instances where symptoms of possible infection preceded symptoms of painful crisis, as suggested by the literature on sickle cell disease. In answer to more detailed queries the program checked the time-order of findings within one document. The potential use of such tables in continuing medical education and other applications in the hospital setting are discussed.

Automatic Application of Health Care Criteria to Narrative Patient Records

**Lynette Hirschman, Naomi Sager,
and Margaret Lyman**

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

Proc. 3rd Ann. Symp. on Comp. Appl. in Med. Care, IEEE, 1979, 105-113.

This paper describes an experimental computer program for the application of health care review criteria to hospital discharge summaries. The use of the computer in this process would make it possible to speed up the routine screening of patient records; it could also facilitate experimental evaluation of alternate proposed audit criteria. The computer program has two components. The first component creates a structured form of the information contained in natural language medical records. It maps the words of each sentence into labelled columns of a table (or *information format*) according to the type of medical information contained in each word. This structured information is suitable for use as a data base in many areas of clinical research. The second component consists of a set of retrieval routines, each of which corresponds to a criterion of the health care evaluation form, e.g., "was the patient afebrile at discharge?" The retrieval component is built up in modular fashion, so that basic routines can be used in other applications. The application of this program to a sample

hospital discharge summary is presented and compared to the results obtained by a physician reviewer.

A CODASYL-Type Schema for Natural Language Medical Records

N. Sager, L. Tick, G. Story, and L. Hirschman

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

Proc. 4th Ann. Symp. on Comp. Appl. in Med. Care, IEEE, 1980, 1027-1033.

This paper describes a CODASYL (network) database schema for information derived from narrative clinical reports. The goal of this work is to create an automated process that accepts natural language documents as input and maps this information into a database of a type managed by existing database management systems. The schema described here represents the medical events and facts identified through the natural language processing. This processing decomposes each narrative into a set of elementary assertions, represented as MEDFACT records in the database. Each assertion in turn consists of a subject and a predicate classed according to a limited number of medical event types, e.g., signs/symptoms, laboratory tests, etc. The subject and predicate are represented by EVENT records which are owned by the MEDFACT record associated with the assertion. The CODASYL-type network structure was found to be suitable for expressing most of the relations needed to represent the natural language information. However, special mechanisms were developed for storing the time relations between EVENT records and for recording connections (such as causality) between certain MEDFACT records. This schema has been implemented using the UNIVAC DMS-1100 DBMS.

Research into Methods for Automatic Classification and Fact Retrieval in Science Subfields

N. Sager, L. Hirschman, C. White, C. Foster, S. Wolff, R. Grad and E. Fitzpatrick

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

Report No. 13, October 1980, 93 pages.

The broad goal of this research was to extend the applicability of techniques for automated language analysis being developed for the processing of information in natural language data stores. The work proceeded on two levels: (1) procedures applicable to English material independent of the subject matter; (2) procedures directed to the special use of language in particular disciplines (the sublanguage of the discipline). Previous work had shown that a computerized

grammar of English could be used in a parsing program to obtain the grammatical structure of input sentences; and subsequent processing, using the word classes special to a given subject area, could arrange the information in the sentences in table-like forms (called information formats). With these methods, textual information which heretofore had only been accessible via key words and other word-scanning methods could be organized in tabular form, similar to numerical and scientific data bases.

To make these methods more widely applicable the problems inherent in word classification have to be faced. Automatic parsing of English text sentences requires a computerized dictionary that gives parts of speech and other grammatical properties of words; mapping parsed sentences into information formats requires a sublanguage dictionary that gives the subject-matter (= semantic) word class membership of the words. The research investigated several computer-based methods for preparing and utilizing such dictionaries.

On the sublanguage level, experiments in automatic word class generation and sublanguage word-class "boot strapping" were undertaken. In the latter, we use patterns of sublanguage word class co-occurrence, known from an initial set of texts, to determine the sublanguage word class of new words in subsequent texts.

Retrieving Time Information from Natural Language Texts

Lynette Hirschman

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

In Information Retrieval Research, Oddy et al. (eds.), London: Butterworths, 1981, 154-171.

An understanding of time relations is central to processing information contained in a narrative. A typical narrative is concerned with the relative ordering or progression of events over time, and the information that one would like to retrieve is often of the type "What happened to event x?" or "Did event x precede event y?" Determination of causality also requires a knowledge of time relations, since an event x can cause an event y only if it precedes event y in time.

There has been considerable interest in the processing of time information among researchers in artificial intelligence. Several systems have been designed which compute time relations from a structured input of time specifications. The work described here differs from this other work in that it focuses specifically on the problem of extracting time information from coherent natural language text.

The program described here was developed in the course of ongoing research in natural language processing at the Linguistic String Project of New York University. This research has been concerned with the creation of a database from natural language input and the retrieval of information from such a database. The work on narrative time was an outgrowth of an experiment on retrieval of information from narrative medical records.

Model-theoretic Pragmatics: Dynamic Models and an Application to Presuppositions and Implicature

Douglas B. Moran

Department of Computer Science
Oregon State University
Corvallis, Oregon 97331

Univ. of Michigan Ph.D. Dissertation, Comp. Studies in Formal Ling. N-22, October 1980, 328 pages.

Model-theoretic semantics is a computationally attractive formalism for the semantics of natural languages. The logical model can be viewed as a database of information about the world with the evaluation of logical formulas (representing the semantics of sentences) retrieving information from this database. However, this formalism has the limitation that the information in a model is complete and static. To overcome this problem, a formalism called *dynamic model-theoretic semantics* is developed and applied to the system given in Montague's *PTQ*. Dynamic models contain incomplete information and can have information added to them as they are used. The evaluation of logical formulas in a dynamic environment produces effects — termed model-theoretic pragmatics — that do not occur when a conventional model is used.

The primary pragmatic effect studied here accounts for presupposition and implicature with *meaning postulates* — logical formulas that must be true for the model to be reasonably used in the given application. With conventional models, meaning postulates are used to select reasonable models; with dynamic models, they are consistency checks on the information being added to the model. Since only new information is checked, this mechanism produces different behaviors for different models (contexts).

Dynamic model-theoretic semantics is non-monotonic: it is impossible to detect when critical elements are missing from the domains of quantifiers or from the domains of functions being compared, and thus the truth-value of a quantifier expression of an equality test may change when the model expands.

A constructed element (a function or set) can have its domain expanded between the time that it is specified (selected) and the time that it is used. A mechanism for maintaining the consistency of an expanded

element with its original specification is presented and shown to have interesting consequences: it permits procedural — as well as declarative — representations of the information in the model, and it permits inconsistent information to be entered. This mechanism also has the capability of providing a motivated restriction on the kinds of inconsistencies that can occur.

Using Semantics in Non-Context-Free Parsing of Montague Grammar

David Scott Warren

Department of Computer Science
State University of New York
Stony Brook, New York 11794

Joyce Friedman

Department of Computer and Communication Sciences
The University of Michigan
Ann Arbor, Michigan 48109

Comp. Studies in Formal Ling. N-27, August 1981, 41 pages.

In natural language processing, questions concerning the appropriate interaction of syntax and semantics have long been of interest. Montague grammar and its fully formalized syntax and semantics provide a complete, well-defined context in which these questions can be considered. This paper describes how semantics can be used during parsing to reduce the combinatoric explosion of syntactic ambiguity in Montague grammar. A parsing algorithm, called semantic equivalence parsing, is described and examples of its operation are given. The algorithm is applicable to general non-context-free grammars that include a formal semantic component. The second portion of the paper places semantic equivalence parsing in the context of the very general definition of an interpreted language as a homomorphism between syntactic and semantic algebras.

Adaptation of Montague Grammar to the Requirements of Parsing

Jan Landsbergen

Philips Research Laboratories
Eindhoven, THE NETHERLANDS

MC Tract 136, Formal Methods in the Study of Language, Mathematics Centre, Amsterdam, 1981, 399-419.

The paper describes a variant of Montague grammar, of which the composition rules have analytical counterparts on which a parsing algorithm can be based. Separate attention is given to the consequences of including rule schemes and syntactic variables in the grammar.

A Framework for Processing Ill-Formed Input**Ralph M. Weischedel**Department of Computer and Information Sciences
University of Delaware
Newmark, Delaware 19711**Norman K. Sondheimer**Software Research MS 2G3
Sperry Univac
Blue Bell, Pennsylvania 19424*Technical Report, 1981, 38 pages.*

If natural language processing systems are ever to pass the Turing test, or if they are ever to achieve natural, cooperative behavior, they must be able to process input that is ill-formed lexically, syntactically, semantically, or pragmatically. Systems must be able to partially understand or at least give specific, appropriate error messages when input does not correspond to their model of language and of context.

Out of our own work and the work of others, we propose meta-rules and a control structure under which they are invoked as a framework for processing ill-formed input. The left-hand-side of a meta-rule diagnoses (hypothesizes) a problem with the input as a violated rule of normal processing. The right-hand-side rewrites a violated rule as a relaxed one and states how processing may be resumed, if at all.

Several specific meta-rules are given as examples. In addition, a sketch is included of how several significant heuristics developed by others can be formulated as meta-rules. An analysis of the limitations of this framework is also provided.

Inference and Control in Multiprocessing Environments**Harold Shubin**Department of Computer Science
State University of New York at Buffalo
Amherst, New York 14226*Technical Report No. 186, September 1981, 56 pages.*

The ideas behind inference, as used in Artificial Intelligence (AI) systems, are similar to those of certain control structures used in other areas of computation. This paper discusses those ideas and specifically studies forward, backward, and bi-directional inference; and the data flow concept, lazy evaluation and bi-directional search. A model of computation called the Supplier-Producer-Consumer (SPC) Model, is introduced as a vehicle for making contrasts between pairs of inference and control strategies. Contrasts along another dimension are made in the model by discussing static and eager evaluation schemes.

Multiprocessing is an idea which has been implemented differently in different areas of computer science. This paper discusses the benefits of software simulations of multiprocessing on uni-processing sys-

tems for these control and inference methods. Finally, bi-directional methods (computation, inference, and search) are discussed. The combination of forward and backward computation methods allows each to assist the other, and suggests new ways for a program to interact with a user.

Analyzing Intention in Dialogues**J.F. Allen and C.R. Perrault**Department of Computer Science
The University of Rochester
Rochester, New York 14627*Report No. TR50, April 1979, 75 pages.*

This paper describes a model of cooperative behavior and shows how such a model can explain some interesting linguistic behavior. We assume that agents attempt to recognize the plans of other agents and then use this plan when deciding what response to make. In particular, we show that, given a setting in which purposeful dialogues occur, this model can account for responses that provide more information than explicitly requested and for appropriate responses to both short sentence fragments and indirect speech acts.

Beyond Question-Answering**P.R. Cohen, C.R. Perrault, and J.F. Allen**Bolt Beranek and Newman Inc.
10 Moulton Street
Cambridge, Massachusetts 02238*Technical Report 4644, May 1981.**(To appear in Strategies for Natural Language Processing, Lehnert and Ringle (eds.), Erlbaum Assoc., in press.)*

We show that users of question-answering systems expect those systems to be responsive to their unstated plans and goals. Techniques needed to accomplish this should be special cases of more general abilities, in particular, the ability to recognize the user's plan and to plan a helpful response. We propose and justify a new system architecture embodying this framework, and we illustrate how that architecture is applied in two implemented systems. The first is a question-answering system, and the second is a simple decision-support system, for which both graphic and linguistic means of communication are available.

Conversational Coherency in Technical Conversations**Rachel Reichman**ISSCO
Université de Genève
17 Rue de Candolle
CH-1205 Genève SWITZERLAND*Working Paper 43, 1979.*

An analysis of technical conversations is presented which uses the framework of *context spaces* earlier developed for the analysis of social exchanges. Both forms of discourse are shown to display similarities which are brought out by this form of analysis. A number of instances of surface linguistic phenomena, such as deictic "that," present progressive tense, pronominalization, and clue words such as "it's like" and "now", are presented and discussed. These hitherto unexplained phenomena are accounted for in terms of the underlying discourse structure and the resulting state and focus level assignments to discourse constituents.

Upward Branching Phrase Markers: The State of the Debate

G. Sampson

ISSCO

Université de Genève

17 Rue de Candolle

CH-1205 Genève SWITZERLAND

Working Paper 45, 1980.

A 1975 paper by the author ("the single mother condition") argued that the definition of *phrase marker* in linguistics should, for reasons both of empirical adequacy and theoretical elegance, be modified to permit both upward as well as downward branching. This paper examines various reactions to this proposal that have been published. Certain criticisms are accepted as valid but not fatal. The rest turn out to be based on a misunderstanding of the original claim.

Three Strategic Goals Employed in Conversational Openings

M. Rosner

ISSCO

Université de Genève

17 Rue de Candolle

CH-1205 Genève SWITZERLAND

Working Paper 46, 1981.

This paper tries to explain a short transcript of a conversational opening as completely as possible within the framework which takes conversational behaviour as defined by the operation of a sophisticated planning mechanism. To account for conversational openings in general and this transcript in particular, it is argued that a crucial role is played by the satisfaction, for each participant, of three strategic goals relating to attention, identification, and greeting. Additional tactics for gaining information are also described as necessary to account for this transcript. The final analysis employs the definitions of the goals and tactics as defined. It is concluded that many of the deficiencies in the final analysis could be avoided by further investigations aimed at formalizing the explanatory framework.

A Computer-Based Teaching Scheme for Creative Writing

Mike Sharples

Department of Artificial Intelligence

University of Edinburgh

Forrest Hill

Edinburgh EH1 2QL SCOTLAND

Proc. 3rd World Conf. on Computers in Education, Lewis and Tagg (eds.), North-Holland, 1981, 483-488.

The paper describes a computer-based teaching scheme for creative writing. The scheme, based on a cognitive theory of the writing process, develops children's meta-linguistic knowledge and applies it to the exploration and improvement of written style. The children use computer programs to generate and transform text and to explore the process of story production. We present an outline description of the scheme and the results of a pilot project with 6 eleven-year-old pupils.

Microcomputers and Creative Writing

Mike Sharples

Department of Artificial Intelligence

University of Edinburgh

Forrest Hill

Edinburgh EH1 2QL SCOTLAND

In Microcomputers in Secondary Education, Howe and Ross (Eds.), London: Kogan Page, 1981, 138-157.

There are few examples of computer aids for education in language arts and, of those few, most provide pupils with drill and practice exercises in grammar, spelling, or writing style. The computer programs described in this paper are different. They offer powerful and general learning aids — a sentence generator, a story planner, a word processor/text transformer, an automated thesaurus and dictionary, a spelling corrector — which may be used both by the teacher to demonstrate language construction and by the pupil to compose and alter text. Being tools rather than teaching systems they are not aligned to a particular syllabus and, as the linguistic information is held in simple data files, the programs may be easily modified to manipulate other languages — from Latin to mathematical expressions. The paper begins with a brief history and critique of computer assisted instruction (CAI) followed by an impression of a computer-based "Workshop" for exploring language and creative writing and, lastly, an account of a project with eleven-year-old children who used a prototype part of the workshop in a creative writing course.