

ABSTRACTS OF CURRENT LITERATURE

In October of 1984, a workshop was held at Duke University to consider problems relating to the design of transportable natural language processors. The workshop was funded in part by a grant from the National Science Foundation and was co-sponsored by Duke and the ACL. A number of papers resulted from the meeting, and several of them appear in the April 1985 issue of the *ACM Transactions on Office Information Systems* (volume 3, number 2). Bruce Ballard guest edited the special issue, and wishes to thank the following persons for their assistance in reviewing papers from the workshop: Joan Bachenko, Ron Brachman, Fred Damerau, Samuel Epstein, Kurt Godden, Ralph Grishman, Carole Hafner, Jerry Hobbs, Karen Jensen, Mark Jones, Karen Kukich, Elaine Marsh, Sharon Salveter, Jonathan Slocum, Norm Sondheimer, Douglas Stumberger, and Craig Thompson. The following are abstracts from that special issue.

Transportable Natural Language Processing through Simplicity – The PRE System

Samuel S. Epstein

Bell Communications Research
435 South Street
Morristown, NJ 07960

PRE (Purposefully Restricted English) is a restricted English database query language whose implementation has addressed engineering goals, namely, habitability, interapplication transportability, performance, and use with a reliable database management system that supports large numbers of concurrent users and large databases. Habitability has not been demonstrated, but initial indications are encouraging. The other goals have clearly been achieved. The existence of the PRE system demonstrates that an explicitly “minimalist” approach to natural language processing can facilitate achievement of transportability.

Transporting the Linguistic String Project System from a Medical to a Navy Domain

Elaine Marsh

Navy Center for Applied Research in
Artificial Intelligence

Carol Friedman

Courant Institute of Mathematical
Sciences

The Linguistic String Project (LSP) natural language processing system has been developed as a domain-independent natural language processing system. Initially utilized for processing sets of medical messages and other texts in the medical domain, it has been used at the Naval Research Laboratory for processing Navy messages about shipboard equipment failures. This paper describes the structure of the LSP system and the features that make it transportable from one domain to another. The processing procedures encourage the isolation of domain-specific information, yet take advantage of the syntactic and semantic similarities between the medical and Navy domains. From our experience in transporting the LSP system, we identify the features that are required for transportable natural language systems.

Portability of Syntax and Semantics in Datalog

Carole D. Hafner, Kurt Godden

General Motors Research Laboratories

This paper presents a discussion of the techniques developed and problems encountered during the design, implementation, and experimental use of a portable natural language processor. Datalog (for “database dialogue”) is an experimental natural language query system, which was designed to achieve a maximum degree of portability and extensibility. Datalog uses a three-level architecture to provide both portability of syntax to new and extended tasks and portability of semantics to new database applications. The implementation of each of the three levels, the structures and conventions that control the interactions among them, and the way in which different aspects of the design contribute to portability are described. Finally, two specific, implemented examples are presented, showing how it was possible to transport or extend Datalog by changing only one “layer” of the system’s knowledge and achieve correct processing of the extended input by the entire system.

Problems and Some Solutions in Customization of Natural Language Database Front Ends

Fred J. Damerau

IBM T. J. Watson Research Laboratory

This paper is concerned with some of the issues arising in the development of a domain-independent English interface to IBM SQL-based program products. The TQA system falls into the class of multilayered natural language processing systems. As a result, there is a large number of potential points at which customization to a particular database can be done. Of these, we discuss procedures that affect the reader, the lexicon, the lowest level of grammar rules, the semantic interpreter, and the output formatter. Our tests lead us to believe that the approach we are taking will make it

ASK Is Transportable in Half a Dozen Ways**Bozena Henisz Thompson,
Frederick B. Thompson**

California Institute of Technology

possible for database administrators to generate robust English interfaces to particular databases without help from linguistic experts.

This paper is a discussion of the technical issues and solutions encountered in making the ASK System transportable. A natural language system can be "transportable" in a number of ways. Although transportability to a new domain is most prominent, other ways are also important if the system is to have viability in the commercial marketplace.

On the one hand, transporting a system to a new domain may start with the system prior to adding any domain of knowledge and extend it to incorporate the new domain. On the other hand, one may wish to add to a system that already has knowledge of one domain the knowledge concerning a second domain, that is, to extend the system to cover this second domain. In the context of ASK, it has been natural to implement extending and then achieve transportability as a special case.

In this paper, we consider six ways in which the ASK System can be extended to include new capabilities:

- to a new domain,
- to a new object type,
- to access data from a foreign database,
- to a new natural language,
- to a new programming language,
- to a new computer family.

Special-purpose applications, such as those to accommodate standard office tasks, would make use of these various means of extension.

Transportability to Other Languages: The Natural Language Processing Project in the AI Program at MCC**Jonathan Slocum, Carol F. Justus**

Microelectronics and Computer Technology Corporation

We discuss a recently launched, long-term project in natural language processing, the primary concern of which is that natural language applications be transportable among human languages. In particular, we seek to develop system tools and linguistic processing techniques that are themselves language-independent to the maximum extent practical. In this paper we discuss our project goals and outline our intended approach, address some cross-linguistic requirements, and then present some new linguistic data that we feel support our approach.

The following abstracts are from *Proceedings of the Conference on Theoretical and Methodological Issues in Machine Translation of Natural Languages*, held at Colgate University, Hamilton, New York, 14-16 August 1985. Copies of the *Proceedings* are no longer available; papers should be requested from the author(s).

A MU₁ View of the <C,A>_T Framework in Eurotra**Doug Arnold, Lieven Jaspaert, Rod Johnson,
Steven Krauwer, Mike Rosner, Louis
des Tombe, Nino Varile, Susan Warwick**
pp. 1-14

The background to this paper is the attempt within EUROTRA to develop a general framework for research and development work in MT, providing in particular an environment which facilitates reasoning about the relationships between the representations that are necessary for automatic translation between natural languages. The more immediate background is the attempt to apply this framework experimentally on a small scale in developing a "proto-EUROTRA" (affectionately, "the toy") over the summer and autumn of this year. The aim of this paper is to give a reasonably clear idea about the user language and theories of representation for this experiment (the Mu₁ level in terms of the paper by Louis des Tombe), and to indicate en route some of the directions for further work. It reports work in progress, and is thus deliberately speculative, programmatic, and rather informal.

For concreteness, section 2 gives a brief and rather casual restatement of some of the key assumptions behind this work.

On the Production Environment Proposed for the Eurotra Project
D. Bachut

Groupe d'Etudes pour la Traduction

We present the general architecture of a production environment which is specific for a M(A)T system, and give some proposals to integrate new functionalities in this system. A good management of the results of the translation process may lead to an easier improvement of the linguistic

Automatique

Université Scientifique et Médicale de
Grenoble
BP 68 - 38402 Saint Martin d'Hères
& Institut de Formation et Conseil en
Informatique
27, rue de Turenne - 3800 Grenoble
FRANCE

pp. 15-26

**A Case Study in Software Evolution:
from Ariane-78.4 to Ariane-85**

**Ch. Boitet, P. Guillaume,
M. Quezel-Ambrunaz**

Groupe d'Etudes pour la Traduction
Automatique
Université Scientifique et Médicale de
Grenoble
BP 68 - 38402 Saint Martin d'Hères,
FRANCE

pp. 27-58

New Approaches to Machine Translation

Jaime G. Carbonell, Masaru Tomita

Carnegie-Mellon University
Pittsburgh, PA 15213

pp. 59-74

Lexicon-Driven Machine Translation

R.E. Cullingford

Georgia Institute of Technology
Atlanta, GA 30332

B.A. Onyshkevych

Princeton University
Princeton, NJ 08544

pp. 75-115

**On the Design of Expert Systems Grafted
on MT Systems**

R. Gerber, Ch. Boitet

Groupe d'Etudes pour la Traduction
Automatique
Université Scientifique et Médicale de
Grenoble

BP 68 - 38402 Saint Martin d'Hères,
FRANCE

pp. 116-134

data.

We describe a possible organisation for the machine environment of such a system and for the management of the data base of texts. Finally, we give some general rules for the implementation of a monitor.

No abstract.

The current resurgence of interest in machine translation is partially attributable to the emergence of a variety of new paradigms, ranging from better translation aids and improved pre- and post-editing methods, to highly interactive approaches and fully automated knowledge-based systems. This paper discusses each basic approach and provides some comparative analysis. It is argued that both interactive and knowledge-based systems offer considerable promise to remedy the deficiencies of the earlier, more ad hoc post-editing approaches.

Machine Translation (MT) systems have historically relied upon explicit grammars in order to analyze the source text and reproduce it in the target language. In this paper, we argue for a style of MT in which the focus of processing is at the level of the *lexicon*, rather than the grammar. This approach to translation allows an analyzer to map source sentences into an interlingual form, which then can be mapped (perhaps after intermediate inferencing steps) back into target sentence(s) which are paraphrase-equivalent to the original. Advantages of the approach include: (1) the possibility for different paraphrases of the original; (2) the capability for multi-sentence expression of the original when no single word (e.g., a verb) exists in the target language which spans the same meaning complex as a word in the source; (3) a uniform approach to word sense disambiguation and anaphoric reference resolution; and, most important, (4) the possibility for robust handling of ungrammatical and ellipsed source text.

Our MT systems integrate many advanced concepts from the fields of computer science, linguistics, and AI: specialized languages for linguistic programming based on production systems, complete linguistic programming environment, multilevel representations, organization of the lexicons around "lexical units", units of translation of the size of several paragraphs, possibility to use text-driven heuristic strategies.

We are now beginning to integrate new techniques: unified design of an "integrated" lexical data base containing the lexical in "natural" and "coded" form, use of the "static grammars" formalism as a specification language, and design of a kind of structural meta-editor (driven by some static grammar) allowing the interactive construction of a document in the same way as syntactic editors are used for developing programs.

This paper centers on our study of possible additions of expert systems equipped with metalinguistic and extralinguistic knowledge, in order to

Machine Translation in the SDCG Formalism

Xiuming Huang

Institute of Linguistics

Chinese Academy of Social Sciences

Beijing, China

mail: Computing Research Laboratory

New Mexico State University

Las Cruces, NM 88003

pp. 135-144

Machine Translation as an Expert Task

Rod Johnson, Pete Whitelock

Centre for Computational Linguistics

University of Manchester Institute of

Science and Technology

P.O. Box 88

Manchester M60 1QD UK

pp. 145-153

The Significance of Sublanguage for Automatic Translation

Richard I. Kittredge

University of Montreal

pp. 154-166

Integrating Syntax and Semantics

Steven L. Lytinen

Cognitive Systems, Inc.

234 Church Street

New Haven, CT 06510

pp. 167-178

solve some problems encountered in second-generation MT systems.

Several examples of the possible use of expert-corrector systems in M(A)T (Machine (Aided) Translation) systems are given.

The paper describes the SDCG (Semantic Definite Clause Grammars), a formalism for Natural Language Processing (NLP), and the XTRA (English Chinese Sentence TRANslator) machine translation (MT) system based on it. The system translates general domain English sentences into grammatical Chinese sentences in a fully automatic manner. It is written in Prolog and implemented on the DEC-10, the GEC, and the SUN workstation.

SDCG is an augmentation of the DCG (Definite Clause Grammar; Pereira et al. 1980) which in turn is based on CFG (Context Free Grammar). Implemented in Prolog, the SDCG is highly suitable for NLP in general, and MT in particular.

A wide range of linguistic phenomena is covered by the XTRA system, including multiple word senses, coordinate constructions, and prepositional phrase attachment, among others.

The case against fully automatic high quality machine translation (FAHQMT) has been well-canvassed in the literature ever since ALPAC. Although considerable progress in computational linguistics has been made since then, many of the major arguments against FAHQMT still hold (a good resumé is given by Martin Kay (1980)).

It is not our intention to reopen the case for FAHQMT here. Rather, we contend that, accepting that FAHQMT is not possible in the current state of the art, it is both feasible and desirable to set up R & D programmes in MT which can both produce results which will satisfy sponsors and provide an environment to support research directed towards bringing MT closer to the ultimate goal of FAHQMT.

This paper describes the rationale and organisation behind one such programme, the UMIST English-Japanese MT project.

This paper address three questions:

- What is sublanguage?
- Why is sublanguage analysis important for automatic translation?
- How can a translation system take advantage of sublanguage properties?

The first of these questions appears to have a simple answer. Natural languages clearly have specialized varieties which are used in reference to restricted subject matter. We speak, for example, of the "language of chemistry" to mean a loosely defined set of sentences or texts dealing with a particular part of reality.

But when we consider the automatic translation of specialized language, we are forced to be more precise. We must describe sublanguages as coherent, rule-based systems. The attempt to write grammars for special-purpose sublanguages raises a number of theoretical and practical problems, which are only now being intensively discussed. But since the only path to high-quality automatic translation seems to lie through sublanguage (at least during the next decade or two), we have no choice but to solve these problems. This paper should therefore be considered as a brief summary and progress report.

Well-known examples such as Bar-Hillel's (1960) "The box is in the pen" illustrate that extensive semantic analysis is necessary to resolve ambiguities that must be resolved in machine translation. If one accepts the premise that semantics should be added to the analysis techniques used in machine translation, what is the way in which it should be added? This paper will argue for an *integrated* approach to semantic processing. That is, syntactic and semantic processing should take place at the same time, rather than in separate stages. However, although I will argue for the inte-

LMT: A PROLOG-Based Machine Translation System

Michael C. McCord

IBM T. J. Watson Research Center
Yorktown Heights, NY 10598

pp. 179-182

Recovering the Speaker's Decisions during Mechanical Translation

David D. McDonald

University of Massachusetts at Amherst,
01003

pp. 183-199

Structural Transformation in the Generation Stage of MU Japanese to English Machine Translation System

Makoto Nagao

Department of Electrical Engineering
Kyoto University
Sakyo, Kyoto, JAPAN

pp. 200-223

Interlingua Design for TRANSLATOR
Sergei Nirenburg, Allen B. Tucker

Department of Computer Science
Colgate University

Victor Raskin

Department of English
Purdue University

pp. 224-244

gration of syntactic and semantic analysis *processes*, I will also argue for the use of a separate body of syntactic *knowledge*, and for building a separate syntactic *representation* during the parsing process. This is in contrast to previous integrated parsers, which have relied almost exclusively on semantic representations to guide the parsing process, and which have not used a separate body of syntactic rules.

The talk will describe a machine translation system, LMT, based in PROLOG, translating from English to German. The effort on LMT per se has just begun this year, although the logic programming methodology for the analysis of the source (English) goes back several years (see, e.g., McCord 1982).

When studied as a source of insight into the human language faculty, rather than to construct a commercially useful service, mechanical translation (MT) is carried out by coupling an otherwise normal natural language parsing system to a normal natural language generation system. In this paper we propose that a crucial capability has been omitted from the design of the parsers that have been used to date, namely a facility for recognizing the information that is implicit in the form of any well written text: matters of emphasis, whether a fact is new or old, whether a relationship is given explicitly or left as an obvious inference, signals of intended moves in the discourse, and other things of this sort. We claim that mechanical translations are "mechanical" principally because they pay no attention to information of this sort, and propose that this can be dealt with by incorporating into the parser knowledge of the relationship between usage and form of the sort that is commonplace in any modern language generation system.

No abstract.

The interlingua approach to machine translation (MT) is characterized by the following two stages:

- 1) translation of the source text into an intermediate representation, an artificial language (*interlingua*) which is designed to capture the various types of meaning of the source text, and
- 2) translation from the interlingua into the target text.

Over the years a number of MT projects tried to develop interlingua-based systems. In these projects the amount of linguistic and encyclopaedic knowledge used to produce intermediate representations was quite limited. However, even at that level difficulties connected with encoding knowledge seemed overwhelming. The TRANSLATOR project at Colgate University benefits from recent developments in knowledge representation techniques. The text of its interlingua text reflects syntactic, lexical, contextual, discourse (including speech situation), and pragmatic meaning of the input. This paper discusses the lexicon and grammar of the interlingua used in TRANSLATOR, and touches upon the structure of the bilingual (source language to interlingua) dictionaries. The actual compilation of the interlingua dictionary and additional knowledge bases is an empirical process during which modifications to the original formulations are expected to occur. At all times in the design process the authors were guided by the desire to make decisions that are "literate" from the point of view of

The Level Hypothesis in Discourse Analysis

James Pustejovsky

Department of Computer and Information Sciences
University of Massachusetts at Amherst,
01003
pp. 245-267

Linguistics and Natural Language Processing

Victor Raskin

Purdue University
pp. 268-282

A Preliminary Linguistic Framework for Eurotra, June 1985

Louis des Tombe, Doug Arnold, Lieven Jaspaert, Rod Johnson, Steven Krauwer, Mike Rosner, Nina Varile, Susan Warwick
pp. 283-288

Feasibility Study of Personal/Interactive Machine Translation Systems

Masaru Tomita

Computer Science Department
Carnegie-Mellon University
Pittsburgh, PA 15213
pp. 289-297

Static Grammars: A Formalism for the Description of Linguistic Models

Bernard Vauquois, Sylviane Chappuy

Groupe d'Etudes pour la Traduction Automatique
Université Scientifique et Médicale de Grenoble
BP 68 - 38402 Saint Martin d'Hères,
FRANCE
& Institut de Formation et Conseil en Informatique
27, rue de Turenne - 3800 Grenoble
FRANCE
pp. 298-322

linguistic theory and the experience of knowledge representation in artificial intelligence.

In this paper I would like to explore some difficult questions related to topics in discourse analysis (henceforth DA) and offer a partial solution to some of them. In particular, I will address the issue of *levels* in DA and how the various approaches taken within the field can be classified according to a leveled model. I then want to consider an approach I have been pursuing for representing the semantics of discourse, and consider how it fits in to the proposed model for DA.

The paper addresses the issue of cooperation between linguistics and natural language processing (NLP), in general, and between linguistics and machine translation (MT), in particular. It focuses on just one direction of such cooperation, namely applications of linguistics to NLP, virtually ignoring for now any possible applications of NLP to linguistics, which can range from providing computer-based research tools and aids to linguistics to implementing formal linguistic theories and verifying linguistic models.

Section 1 deals with the question why linguistics must be applied to NLP and what the consequences of ignoring it are. Section 2 provides a counterpoint by discussing how linguistics should not be applied to NLP and, by contrast and inference, how it should be. Section 3 narrows the discussion down to one promising approach to NLP, the sublanguage deal, and the interesting ways in which linguistics can be utilized within a limited sublanguage. Section 4 is devoted specifically – but necessarily briefly – to the things linguistics can contribute to MT.

The work described here was the consequence of the idea that we wanted to make a new, more interesting theoretical start in EUROTRA. It is preliminary and not fully developed yet; it should be seen as the reflection of a way of thinking about MT. Currently, we are making it more precise, and experimenting with it. In this paper, we sketch the general outlines of the new EUROTRA framework; some more exemplification can be found in the paper by D.J. Arnold et al. (p. 1 of this volume).

Most existing practical machine translation systems are designed to translate documentation, such as technical papers and manuals. However, there is a growing need for translating not only large texts but also personal short texts such as letters and informal messages. The conventional machine translation systems, which are intended to translate large texts, are not very suitable for these kinds of small jobs. We need an interactive system which has a totally different design philosophy. This paper describes the design philosophy of personal/interactive machine translation systems, and studies in feasibility.

For a linguistic model it is necessary, first of all, to define the mapping between the strings of words of a language and their structural organisation, given that with transducers there are many ways of obtaining the same result using different strategies.

This mapping, which we will call a “static grammar”, is independent of the analysis, generation, or whatever strategy adopted. Moreover, the formalism of a static grammar is not affected by the choice or number of interpretation levels.

Such a grammar is the “reference” for any dynamic modular rule organisation, whether analysis or generation.

We present here a “static grammar” formalism recently developed at Grenoble (G.E.T.A. Groupe d'Etude pour la Traduction Automatique) under the supervision of Prof. B. Vauquois.

Using this formalism, any given language can be described as a series

of "charts". Each "chart" describes how a certain group of strings corresponds to the structure associated with this group of strings (this structure is a valid and complete substructure of the linguistic model). The structures of all the sentences of a language for a given linguistic model can be described by means of a series of chart inter-references.

The static grammar is used as a base for writing dynamic analysis and generation modules, however, the static grammar does not concern itself with strategic, combinatorial, ambiguity problems or the chart of structures related to dynamic grammars.

We will present here several examples of charts and discuss the dynamic uses of these static grammars.

On Debugging Environment Proposed for Eurotra

N. Verastegui

Institut de Formation et Conseil en
Informatique
27, rue de Turenne - 3800 Grenoble
FRANCE

pp. 323-334

A proposal of external specification of the user environment for the EUROTRA project is presented. The needs of the users and the functions which are necessary for an efficient testing environment are analyzed.

Knowledge Resource Tools for Accessing Large Text Files

Donald E. Walker

Artificial Intelligence and Information
Science Research
Bell Communications Research
435 South Street MRE 2A379
Morristown, NJ 07960

pp. 335-347

This paper provides an overview of a research program just being defined at Bellcore. The objective is to develop facilities for working with large document collections that provide more refined access to the information contained in the "source" materials than is possible through current information retrieval procedures. The tools being used for this purpose are machine-readable dictionaries, encyclopedias, and related "resources" that provide geographical, biographical, and other kinds of specialized knowledge. A major feature of the research program is the exploitation of the reciprocal relationships between *sources* and *resources*. These interactions between texts and tools are intended to support experts who organize and use information in a workstation environment. Two systems under development will be described to illustrate the approach: one providing capabilities for full-text subject assessment; the other for concept elaboration while reading text. Progress in the research depends critically on developments in artificial intelligence, computational linguistics, and information science to provide a scientific base, and on software engineering, database management, and distributed systems to provide the technology.

Reflections on the Knowledge Needed to Process Ill-Formed Language

Ralph M. Weischedel, Lance A. Ramshaw

Bolt Beranek and Newman Inc.
10 Moulton Street
Cambridge, MA 02238

pp. 348-358

This paper reflects about the kinds of morphological syntactic, semantic, and pragmatic knowledge needed to process ill-formed input. We conclude that an excellent start on processing ill-formed input has been exemplified in a number of concrete implementations, but that a substantial amount of fundamental work must still be done if our systems are to understand language robustly to the degree that humans do. Furthermore, we conclude that studying ill-formed language offers important perspectives on the knowledge and architecture needed to correctly understand natural languages.

Characteristics of the METAL Machine Translation System at Production Stage

John S. White

Siemens Communication Systems

pp. 359-369

The METAL machine translation system, a joint project of the Linguistic Research Center and Siemens, has been released for use as part of marketed translation systems. The system, which presently translates technical German into English, is an outgrowth of a traditional, generative approach to automatic analysis and synthesis of natural language phenomena carried out at the Linguistics Research Center for many years. In its present manifestation, it is a modular design consisting of purely monolingual lexicons, transfer lexicons, and an augmented phrase structure grammar. The grammar is powerful enough to constrain application, to build new nodes with essential characteristics of their sons and new synthetic information as well, and to perform transformations to re-order, delete, and create

Relevance, Points of View and Dialogue Modelling

Yorick Wilks

Computing Research Laboratory
New Mexico State University
Las Cruces, NM 88003
pp. 370-387

constituents. The parser is enhanced to allow application of rules in levels, and eliminating unlikely paths via preferential weightings calculated from lexical and grammatical data. The METAL system, conceived in recent years as destined for implementation, has an orientation to user interface which includes sophisticated text stripping, unfound word handling and reconstitution, and a convenient means of working with the lexicons interactively.

This paper attempts to compare two approaches to the modelling of human discourse and, more particularly, dialogue. Both place themselves within a general "information processing paradigm", and both descend from the insights of Grice (1975) that understanding is a matter of inference from what is said and what is assumed. So general is that assumption now, and so widespread are the disciplines that draw upon it – philosophy, psychology, linguistics, and artificial intelligence – that it is hard to capture briefly except in opposition to the transformational-generative paradigm of language, with its notions of the primacy and autonomy of syntax, and the theoretical primacy of explications of competence over those of performance. The Generative Semanticists attempted to merge the two traditions and their failure has made it easier to separate off and clarify the work under discussion here.

The following abstracts are from the *Proceedings of the Ninth International Joint Conference on Artificial Intelligence (IJCAI85)*, Los Angeles, California, August 1985. Two volumes; edited by Aravind Joshi. [ISBN 0-934613-02-8, \$40.00 for AAAI members; \$55.00 for non-members. Morgan Kaufmann Publishers, Inc.; 95 First Street; Los Altos, CA 94022 (415 941-4960).]

Computational Neurolinguistics – What Is It All About?

Helen M. Giggley

Department of Computer Science
University of New Hampshire
pp. 260-266

Computational Neurolinguistics (CN) integrates artificial intelligence (AI) methods with concepts of neurally motivated processing to develop cognitive models of natural language processing.

HOPE is one example of a model developed to address issues in CN. The model is parallel, and exemplifies language as the result of *time* synchronized processes which are asynchronous in nature. Furthermore, the model is substantially validated to include normal behavioral evidence in its design. In addition, it attends to aspects of language breakdown which are well documented in the literature of neurolinguistics or aphasia.

This paper discusses assumptions which underlie the CN approach to model development. It will describe the neurally motivated or "natural computational" processes which produce the model's observable and verifiable behavioral results. The differences in the CN approach to other models of parallel memory process and behavior will be presented. Finally, the contribution of the CN research approach as a tool for investigating the breakdown of language performance and its potential contribution to understanding brain function will be discussed.

A Short Note on Opportunistic Planning and Memory in Arguments

Lawrence Birnbaum

Department of Computer Science
Yale University
New Haven, CT 06520
pp. 281-283

Engaging in an argument is a complex task of natural language processing that involves understanding an opponent's utterances, discovering what his "point" is, determining whether his claims are believable, and fashioning a coherent rebuttal. Accomplishing these tasks requires the coordination of many different abilities and many different kinds of knowledge. Because arguing, and conversation generally, involve real-time interaction with another agent, this coordination must be even more flexible than is required for other natural language processing tasks. An arguer must have some expectations about what his opponents might say, but must also be able to respond to the unexpected. He must have some idea of the claims he wants to make, and plans for putting them forward, but his opponent may confound these plans. Or, more positively, his opponent may say something that offers an unforeseen opportunity to make a point. Arguing

A Process Model of Case-Based Reasoning in Problem Solving

**Janet L. Kolodner, Robert L. Simpson, Jr.,
Katia Sycara-Cyranski**

School of Information and Computer Science
Georgia Institute of Technology
Atlanta, GA 30332
pp. 289-290

Learning to Understand Contractual Situations

**Seth R. Goldman, Michael G. Dyer,
Margot Flowers**

Artificial Intelligence Laboratory
Computer Science Department
University of California
Los Angeles, CA 90024
pp. 291-293

Granularity

Jerry R. Hobbs

Artificial Intelligence Center
SRI International
Menlo Park, CA 94025
& Center for the Study of Language and Information
Stanford University
Stanford, CA 94305
pp. 432-435

Naive Kinematics: One Aspect of Shape

Yoav Shoham
Computer Science Department
Yale University
New Haven, CT 06520
pp. 436-442

A Representation for Complex Physical Domains

thus exemplifies the need for a flexible mix of top-down and bottom-up processing in both language understanding and production.

This paper is concerned with the roles of memory processing and planning in the processes of understanding and generating utterances in an argument or conversation. In particular, I will show that the memory and inferential processing necessary in order to understand another person's utterances can and should perform much of the work required to generate a response, work that most previous theories of conversation would delegate to explicit, goal-directed planning. The consequences of this, both for memory processing and for planning, will be briefly described and analyzed.

Much of the problem solving done by both novices and experts uses "case-based" reasoning, or reasoning by analogy to previous similar cases. We explore the ways in which case-based reasoning can help in problem solving. According to our model, transfer of knowledge between cases is guided largely by the problem solving process itself. Our model shows the interactions between problem solving processes and memory for experience. Our computer program, called the MEDIATOR, illustrates case-based reasoning in interpreting and resolving common sense disputes.

In the field of law, decisions in previous cases often play a significant role in the presentation and outcome of new cases. Lawyers are constantly recalling old cases to aid them in preparing their own briefs. How do lawyers remember cases? What are the features they use to organize and retrieve past decisions? How do lawyers learn which features are important? To address these questions we are constructing a model of legal novices (i.e., first year law students) and the processes by which they learn contract law. Our model is embodied in a computer program called STARE (from the latin, *stare decisis*, which refers to the principle of using past cases to decide current disputes). STARE will read descriptions of contractual situations and attempt to predict the decision based on its general commonsense knowledge of agreements, the previous cases stored in an episodic memory, and knowledge of some basic legal concepts.

This paper presents a framework for a theory of granularity which is seen as a means of constructing simple theories out of more complex ones. A transitive indistinguishability relation can be defined by means of a set of relevant predicates, allowing simplification of a theory of complex phenomena into computationally tractable local theories, or granularities. Nontransitive indistinguishability relations can be characterized in terms of relevant partial predicates, and idealization allows simplification into tractable local theories. Various local theories must be linked with each other by means of articulation axioms, to allow shifts of perspective. Such a treatment of granularity must be built into the very foundations of the reasoning processes of intelligent agents in a complex world.

Ways in which physical objects interact are explored, and in particular the concept of *freedom* is analyzed. Intuitively, the *fit* between two shapes in a given spatial configuration is a statement about how much one shape needs to be mutilated in order to be made identical to the other. The *freedom* of one object with respect to another specifies what motions the first object can go through without the second one moving. The formulations, termed *naive kinematics*, are compared to work that was done in the kinematics of machinery in the 19th century and that has since been somewhat neglected.

We are exploring a system, called PROMPT, that will be capable of reasoning from first principles and high level knowledge in complex, physical

Sanjaya Addanki

IBM T.J. Watson Research Center
Ernest Davis
 New York University
 pp. 443-446

ONYX: An Architecture for Planning in Uncertain Environments

**Curtis Langlotz, Lawrence Fagan,
 Samson Tu, John Williams**
 Medical Computer Science Group
Branimir Sikic

Department of Medicine
 Knowledge Systems Laboratory
 Stanford University
 Stanford, CA 94305
 pp. 447-449

Understanding Behavior Using Consolidation

Tom Bylander, B. Chandrasekaran
 Laboratory for Artificial Intelligence
 Research

Department of Computer and Informa-
 tion Science
 The Ohio State University
 Columbus, OH 43210
 pp. 450-454

A Decidable First-Order Logic for Knowledge Representation

Peter F. Patel-Schneider

Schlumberger Palo Alto Research
 3340 Hillview Avenue
 Palo Alto, CA 94304
 pp. 455-458

Two Results on Default Logic

Witold Lukaszewicz

Institute of Informatics
 University of Warsaw
 P.O. Box 1210
 00-901 Warszawa, Poland
 pp. 459-461

On the Descriptive Complexity of Production Systems

Peter Trum

Battelle-Institute e.V.

domains. Such problem solving calls for a representation that will support the different analyses techniques required (e.g., differential, asymptotic, perturbation, etc.). Efficiency considerations requires that the representation also support heuristic control of reasoning techniques. This paper lays the ground work for our effort by briefly describing the ontology and the representation scheme of PROMPT. Our ontology allows reasoning about multiple pasts and different happenings in the same space-time. The ontology provides important distinctions between materials, objects, bulk and distributed abstractions among physical entities. We organise world knowledge into "prototypes" that are used to focus the reasoning process. Problem solving involves reasoning with and modifying prototypes.

The ONYX program is designed to fill the need for planning in application areas where traditional planning methodology is difficult to apply. While the program being developed will assist with the planning of cancer therapy, its architecture is intended to be of use whenever goals are ill-specified, plan operators have uncertain effects, or trade-offs and unresolvable conflicts occur between goals. We describe a planning process which uses strategic information and a mechanistic model of the domain. The process consists of three steps: (1) generate a small set of plausible plans based on current data, (2) simulate those plans to predict their possible consequences, and (3) based on the results of those simulations, rank the plans according to how well each meets the goals for the situation.

In this paper we wish to make three contributions to Naive Physics in the context of reasoning about devices. (1) We discuss some limitations of current qualitative simulation approaches with regard to a number of issues in understanding device behavior and point to the need for additional processes. (2) We introduce a new approach to deriving the behavior of devices called consolidation. In this approach, the behavior of a device is derived from the behavior of its components by inferring the behavior of selected substructures of the device. (3) We present an ontology of behavior and structure which is well-suited to the consolidation process. This ontology makes it possible to state rules of behavior composition, i.e., simple patterns of behavior and structure are used to infer additional behaviors.

Even though logic has played an important role in knowledge representation (KR) research, there has been little effort expended on devising decidable logics for KR. Most modifications to logic suggested for KR are either extensions to first-order logic (e.g., to handle non-monotonicity) or ad hoc changes in its inference mechanism. This paper presents a variant of first-order relevance logic that has a decidable algorithm for determining tautological entailment. Although this logic is considerably weaker than standard first-order logic, it can be used effectively in a KR system when semantically correct answers to queries are required within a finite amount of time.

We focus on default logic, a formalism introduced by Reiter to model default reasoning. The paper consists of two parts. In the first one a translation method of non-normal defaults into the normal ones is given. Although not generally valid, this translation seems to work for a wide class of defaults. In the second part a semantics for normal default theories is given and the completeness theorem is proved.

In this paper we formalize different methods for describing control knowledge in production systems by the concept of production system schemes. In this framework these methods are compared from the viewpoint of descriptive complexity, giving us more insight by which means problems

Am Roemerhof 35
D-6000 Frankfurt am Main
West Germany
pp. 462-464

Evidential Reasoning in Semantic Networks: A Formal Theory
Lokendra Shastri, Jerome A. Feldman
Computer Science Department
University of Rochester
Rochester, NY 14627
pp. 465-474

An Endorsement-Based Plan Recognition Program
Michael Sullivan, Paul R. Cohen
Department of Computer and Information Science
University of Massachusetts
Amherst, MA 01003
pp. 475-479

A Guide to the Modal Logics of Knowledge and Belief: Preliminary Draft
John Y. Halpern
IBM Research Laboratory
San Jose, CA 95193
Yoram Moses
Computer Science Department
Stanford University
Stanford, CA 94305
& IBM Research Lab, San Jose
pp. 480-490

Belief, Awareness, and Limited Reasoning: Preliminary Report
Ronald Fagin, Joseph Y. Halpern
IBM Research Laboratory
San Jose CA 95193
pp. 491-501

can be described in an easy and succinct way.

This paper presents an evidential approach to knowledge representation and inference wherein the principle of maximum entropy is applied to deal with uncertainty and incompleteness. It focuses on a restricted representation language – similar in expressive power to semantic network formalisms, and develops a formal theory of evidential inheritance within this language. The theory applies to a limited, but we think interesting, class of inheritance problems including those that involve exceptions and multiple inheritance hierarchies. The language and the accompanying evidential inference structure provide a natural treatment of defaults and conflicting information. The evidence combination rule proposed in this paper is incremental, commutative and associative and hence shares most of the attractive features of the Dempster-Shafer evidence combination rule. Furthermore, it is demonstrably better than the Dempster-Shafer rule in the context of the problems addressed in this paper. The resulting theory can be implemented as a highly parallel (connectionist) network made up of active elements that can solve inheritance problems in time proportional to the depth of the conceptual hierarchy.

This paper describes a plan-recognition program built to explore the theory of endorsements (Cohen 1983). The program evaluates alternative interpretations of user actions and reasons about which are the most likely explanation of the user's intentions. Uncertainty about the various alternatives was encoded in data structures called *endorsements*. The paper describes the workings of this program and the successes and limitations of the endorsement-based approach.

We review and re-examine possible-worlds semantics for propositional logic of knowledge and belief with four particular points of emphasis: (1) we show how general techniques for finding decision procedures and complete axiomatizations apply to models for knowledge and belief, (2) we show how sensitive the difficulty of the decision procedure is to such issues as the choice of modal operators and the axiom system, (3) we discuss how notions of common knowledge and implicit knowledge among a group of agents fit into the possible-worlds framework, and (4) we consider to what extent the possible-worlds approach is a viable one for modelling knowledge and belief. As far as complexity is concerned, we show among other results that while the problem of deciding satisfiability of an S5 formula with one knower is *NP*-complete, the problem for many knowers is *PSPACE*-complete. Adding an implicit knowledge operator does not change the complexity substantially, but once a common knowledge operator is added to the language, the problem becomes complete for exponential time.

Several new logics for belief and knowledge are introduced and studied, all of which have the property that agents are not logically omniscient. In particular, in these logics, the set of beliefs of an agent does not necessarily contain all valid formulas. Thus, these logics are more suitable than traditional logics for modelling beliefs of humans (or machines) with limited reasoning capabilities. Our first logic is essentially an extension of Levesque's logic of implicit and explicit belief, where we intend to allow multiple agents and higher-level belief (i.e., beliefs about beliefs). Our second logic deals explicitly with "awareness", where, roughly speaking, it is necessary to be aware of a concept before one can have beliefs about it. Our third logic gives a model of "local reasoning", where an agent is

**A Computational Theory of Belief
Introspection**

Kurt Konolige

Artificial Intelligence Center
SRI International
Menlo Park, CA 94025
pp. 502-508

**A Model-Theoretic Analysis of Monotonic
Knowledge**

Moshe Y. Vardi

CSLI, Ventura Hall
Stanford University
Stanford, CA 94305
pp. 509-512

**Using Situation Descriptions and Russellian
Attitudes for Representing Beliefs and
Wants**

Alfred Kobsa

Austrian Research Institute for Artificial
Intelligence
Schottengasse 3
A-1010 Vienna, Austria
& Sonderforschungsbereich
Department of Computer Science
Universität des Saarlandes
D-66 Saarbrücken, W. Germany
pp. 513-515

A Procedural Logic

**Michael P. Georgeff, Amy L. Lansky,
Pierre Bessiere**

Artificial Intelligence Center
SRI International
Menlo Park, CA 94025
pp. 516-523

Event Calculus

Gary C. Borchardt

Coordinated Science Laboratory
University of Illinois
Urbana, IL 61801
pp. 524-527

A Common-Sense Theory of Time

viewed as a “society of minds”, each with its own cluster of beliefs, which may contradict each other.

Introspection is a general term covering the ability of an agent to reflect upon the workings of his own cognitive functions. In this paper we will be concerned with developing an explanatory theory of a particular type of introspection: a robot agent’s knowledge of his own beliefs. The development is both descriptive, in the sense of being able to capture introspective behavior as it exists; and prescriptive, in yielding an effective means of adding introspective reasoning abilities to robot agents.

We present a semantic model for knowledge with the following properties: (1) Knowledge is necessarily correct, (2) agents are logically omniscient, i.e., they know all the consequences of their knowledge, and (3) agents are positively introspective, i.e., they are aware of their knowledge, but not negatively introspective, i.e., they may not be aware of their ignorance. We argue that this is the appropriate model for implicit knowledge. We investigate the properties of the model, and use it to formalize the notion of circumscribed knowledge.

A representation scheme for arbitrary beliefs and wants of an agent in respect to a situation, as well as to arbitrary beliefs and wants of other agents, is presented. The representation makes use of elementary situation descriptions (which are formulated in KL-ONE and delimited by partitions), and acceptance attitudes in respect to these descriptions, or to the attitudes thereabout. The scheme forms the representational base of VIE-DPM, the user modelling component of the German-language dialogue system VIE-LANG.

Much of our commonsense knowledge about the real world is concerned with the way things are done. This knowledge is often in the form of *procedures* or *sequences* of actions for achieving particular goals. In this paper, a formalism is presented for representing such knowledge based on the action of *process*. A declarative semantics for the representation is given, which allows a user to state *facts* about the effect of doing things in the problem domain of interest. An operational semantics is also provided, which shows *how* this knowledge can be used to achieve given goals or to form intentions regarding their achievement. The formalism also serves as an executable program specification language suitable for constructing complex systems.

This paper presents Event Calculus, a model for representing the identifying characteristics of physical events in terms of changes in a scene and time-related combinations of other physical events. The model is used to construct a knowledge-based system for event recognition which forms a high-level description of changes in a scene, given a low-level description as input.

Time-varying information is represented in the form of “GRAPHS”, data structures which plot the elements of various domains against time. Several varieties of operations are presented which map GRAPHS into GRAPHS, and representations of physical events are formed as symbolic expressions involving these operations. The paper concludes with an overview of the event recognition system, as implemented in INTERLISP on a VAX 11/780, and an example of a session with this system.

The literature on the nature and representation of time is full of disputes

James F. Allen, Patrick J. Hayes

Departments of Computer Science and
Philosophy
University of Rochester
Rochester, NY 14627
pp. 528-531

**An Essential Hybrid Reasoning System:
Knowledge and Symbol Level Ac-
counts of KRYPTON****Ronald J. Brachman**

AT&T Bell Laboratories
600 Mountain Avenue
Murray Hill, NJ 07974

Victoria Pigman Gilbert

Schlumberger Palo Alto Research
3340 Hillview Avenue
Palo Alto, CA 94304

Hector J. Levesque

Department of Computer Science
University of Toronto
Toronto, Ont., Canada M5S 1A7
pp. 532-539

**The Layered Architecture of a System for
Reasoning about Programs****Charles Rich**

The Artificial Intelligence Laboratory
Massachusetts Institute of Technology
Cambridge, MA 02139
pp. 540-546

**The Restricted Language Architecture of a
Hybrid Representation System****Marc Vilain**

BBN Laboratories
10 Moulton Street
Cambridge, MA 02238
pp. 547-551

**Proportionality Graphs, Units Analysis, and
Domain Constraints: Improving the Power
and Efficiency of the Scientific Discovery
Process****Brian Falkenhainer**

Department of Computer Science
University of Illinois
1304 W. Springfield Avenue
Urbana, IL 61801
pp. 552-554

and contradictory theories. This is surprising since the nature of time does not cause any worry for people in their everyday coping with the world. What this suggests is that there is some form of common sense knowledge about time that is rich enough to enable people to deal with the world, and which is universal enough to enable cooperation and communication between people. In this paper, we propose such a theory and defend it in two ways. We axiomatize a theory of time in terms of intervals and the single relation MEET. We then show that this axiomatization subsumes Allen's interval-based theory. We then extend the theory by formally defining the beginnings and endings of intervals and show that these have the properties we normally would associate with points. We distinguish between these point-like objects and the concept of *moment* as hypothesized in discrete time models. Finally, we examine the theory in terms of each of several different models.

Hybrid inference systems are an important way to address the fact that intelligent systems have multifaceted representational and reasoning competence. KRYPTON is an experimental prototype that completely handles both terminological and assertional knowledge; these two kinds of information are tightly linked by having sentences in an assertional component be formed using structured complex predicates defined in a complementary terminological component. KRYPTON is unique in that it combines in a completely integrated fashion a frame-based description language and a first-order resolution theorem-prover. We give here both a formal Knowledge Level view of the user interface to KRYPTON and the technical Symbol Level details of the integration of the two disparate components, thus providing an essential picture of the abstract function that KRYPTON computes and the implementation technology needed to make it work. We also illustrate the kind of complex question the system can answer.

Cake is a hybrid system which provides reasoning facilities for the Programmer's Apprentice. This paper describes the architecture of Cake, which is divided into eight layers, each with associated representations and reasoning procedures. The operation of Cake is illustrated by a complete trace of the solution of an example reasoning problem. We also argue that a hybrid system in general is characterized by the use of multiple representations in the sense of multiple data abstractions, which does not necessarily imply distinct implementation data structures.

Hybrid architectures have been used in several recent knowledge representation systems. This paper explores some distinctions between various hybrid representation architectures, focusing in particular on systems built around restricted representation languages. This restricted language architecture is illustrated by describing KL-TWO, a hybrid reasoner based on the restricted representation facility RUP. The bulk of this paper discusses KL-TWO, its subcomponents, and the techniques used to interface them.

An important subproblem of scientific discovery is *quantitative discovery*, finding formulas that relate some set (or subset) of a collection of numerical parameters. Current work in quantitative discovery suffers from a lack of efficiency and generality. This paper discusses methods that are efficient and yet general for discovery equations which try to avoid exponential search. Importantly, these methods can derive equations that cover subsets of the data and derive explicit descriptions of when the equations are applicable. These methods are fully implemented in a system named ABACUS, which is described and some of its results are presented.

**A New Kind of Finite-State Automaton:
Register Vector Grammar**

Glenn David Blank

Lehigh University
CSEE Department
Packard Lab 19
Bethlehem, PA 18015
pp. 749-755

Register Vector Grammar is a new kind of finite-state automaton that is sensitive to context – without, of course, being context-sensitive in the sense of Chomsky hierarchy. Traditional automata are functionally simple: symbols match by identity and change by replacement. RVG is functionally complex: ternary feature vectors (e.g., $+ - \pm + - \pm +$) match and change by masking (\pm matches but does not change any value). Functional complexity – as opposed to the computational complexity of non-finite memory – is well suited for modelling multiple and discontinuous constraints. RVG is thus very good at handling the permutations and dependencies of syntax (wh-questions are explored as an example). Because center-embedding in natural languages is in fact very shallow and constrained, context-free power is not needed. RVG can thus be guaranteed to run in a small linear time, because it is FS, and yet can capture generations and constraints that functionally simple FS grammars cannot.

**An Efficient Context-Free Parsing
Algorithm for Natural Languages**

Masaru Tomita

Computer Science Department
Carnegie-Mellon University
Pittsburgh, PA 15213
pp. 756-764

This paper introduces an efficient context-free parsing algorithm and emphasizes its practical value in natural language processing. The algorithm can be viewed as an extended LR parsing algorithm which embodies the concept of a “graph-structured stack”. Unlike the standard LR, the algorithm is capable of handling arbitrary non-cyclic context-free grammars including ambiguous grammars, while most of the LR parsing efficiency is preserved. The algorithm seems more efficient than any existing algorithms including the Docke-Younger-Kasami algorithm and Earley’s algorithm, as far as practical natural language is concerned, due to utilization of LR parsing tables. The algorithm is an all-path parsing algorithm; it produces all possible parse trees (a parse forest) in an efficient representation called a “shared-packed-forest”. This paper also shows that Earley’s forest representation has a defect and his algorithm cannot be used in natural language processing as an all-path parsing algorithm.

Unrestricted Gapping Grammars

Fred Popowich

Natural Language Group
Laboratory for Computer and Communications Research
Computing Science Department
Simon Fraser University
Burnaby, B.C., Canada V5A 1S6
pp. 765-768

Since the introduction of *metamorphosis grammars* (MGs) (Colmerauer, 1978), with their associated *type O*-like grammar rules, there has been a desire to allow more general rule formats in logic grammars. *Gaps*, which refer to strings of unspecified symbols, were added to the MG rule, resulting in *extraposition grammars* (XGs) (Pereira 1981) and *gapping grammars* (GGs) (Dahl and Abramson 1984). *Unrestricted gapping grammars*, which provide an even more general rule format, possess rules of the form “ $\alpha \rightarrow \beta$ ” where α and β may contain any number of terminal, non-terminal, or gap symbols in any order. FIGG, a Flexible Implementation of Gapping Grammars, is an implementation of a large subset of unrestricted GGs which allows either bottom-up or top-down parsing of sentences. This system provides more built-in control facilities than previous logic grammar implementations, which allows the user to restrict the applicability of the rules, and to create grammar rules that will be executed more efficiently.

Parsing with Assertion Sets and Information Monotonicity

G. Edward Barton, Jr., Robert C. Berwick

MIT Artificial Intelligence Laboratory
545 Technology Square
Cambridge, MA 02139
pp. 769-771

We propose a new approach to parsing ambiguity in which a parser always moves forward with the common elements of competing syntactic analyses. The approach involves *assertion sets* constrained so that information is monotonically preserved throughout a parse. Assertion sets have several advantages over trees as a parsing representation. They may also lead to better computational understanding of the attention-shifting mechanism.

Weighted Interaction of Syntax and Semantics in Natural Language Analysis

Leonard Lesmo, Pietro Torasso

Dipartimento di Informatica
Università di Torino

The present paper discusses the extensions to the parsing strategies adopted for FIDO (a Flexible Interface for Database Operations). The parser is able to deal with ill-formed inputs (syntactically ill-formed sentences, fragments, conjunctions, etc.) because of the strict cooperation among syntax and semantics. The syntactic knowledge is represented by

Via Valperga Caluso, 37
10125 Torino - Italy
pp. 772-778

Syntax, Preference and Right Attachment
Yorick Wilks, Xiuming Huang, Dan Fass

Computing Research Laboratory
New Mexico State University
Las Cruces, NM 88003
pp. 779-784

Controlling Search in Flexible Parsing
Steven Minton, Philip J. Hayes, Jill Fain

Computer Science Department
Carnegie-Mellon University
Pittsburgh, PA 15213
pp. 785-787

**Grammatical Relations as the Basis for
Natural Language Parsing and Text
Understanding**

*Samual Bayer, Leonard Joseph,
Condace Kalish*

A045, Mitre Corporation
Bedford, MA 01730
pp. 788-790

**The Role of Perspective in Responding to
Property Misconceptions**

Kathleen F. McCoy

Department of Computer and Informa-
tion Science
University of Pennsylvania
Philadelphia, PA 19104
pp. 791-793

Tailoring Explanations for the User
Kathleen R. McKeown, Kevin Matthews

Department of Computer Science
Columbia University
New York, NY 10027
Myron Wish
AT&T Bell Laboratories
600 Mountain Avenue

means of packets of condition-action rules associated with syntactic categories. The non-determinism is mainly handled by means of rules which restructure the parse tree (called "natural changes") so that the use of backtracking is strongly limited.

In order to deal with difficult cases in which no clear-cut mechanism exists for excluding an interpretation, a weighting mechanism has been added to the parser so that it is possible to explore a few different hypotheses in parallel and to choose the best one on the basis of complex interaction among syntax and semantics.

The paper claims that the right attachment rules for phrases originally suggested by Frazier and Fodor are wrong, and that none of the subsequent patchings of the rules by syntactic methods have improved the situation. For each rule there are perfectly straightforward and indefinitely large classes of simple counterexamples. We then examine suggestions by Ford et al., Schubert and Hirst which are quasi-semantic in nature and which we consider ingenious but unsatisfactory. We offer a straightforward solution within the framework of preference semantics, and argue that the principal issue is not the type and nature of information required to get appropriate phrase attachments, but the issue of where to store the information and with what processes to apply it. We present a Prolog implementation of a best first algorithm covering the data and contrast it with closely related ones, all of which are based on the preferences of nouns and prepositions, as well as verbs.

Most natural language parsers require their input to be grammatical. This significantly constrains the search space that they must explore during parsing. Parsers which attempt to recover from extragrammatical input contend with a search space that is potentially much larger, since they cannot necessarily prune branches when grammatical expectations are violated. In this paper we discuss the control structure of the experimental MULTIPAR parser, which directs its search by exploring potential parses in order of their degree of grammatical deviation.

The KING KONG parser described by this paper attempts to apply the principles of relational grammar to the parsing of English in order to overcome the problems encountered by syntactic and semantic parsers. Specifically, this parser uses relational categories such as subject, direct object, and instrument to map syntactic constituents onto semantic roles within CD-like structures. Thus, the parser makes use of both syntactic and semantic information to guide its parse.

In order to adequately respond to misconceptions involving an object's properties, we must have a context-sensitive method for determining object similarity. Such a method is introduced here. Some of the necessary contextual information is captured by a new notion of *object perspective*. It is shown how object perspective can be used to account for different responses to a given misconception in different contexts.

In order for an expert system to provide the most effective explanations, it should be able to tailor its responses to the concerns of the user. One way in which explanations may be tailored is by point of view. A method is presented for representing the knowledge to support different points of view in the current domain. In addition, we present a method for determining the point of view to take by inferring the user's goal within a brief discourse segment. The advising system's response to the derived goal depends on the strength of its belief in the inference for which a method of

Murray Hill, NJ 07974
pp. 794-798

Description-Directed Natural Language Generation

David D. McDonald, James D. Pustejovsky

Department of Computer and Information Science

University of Massachusetts at Amherst,
01003

pp. 799-805

determination is also provided. This information enables the system to decide what answer to give to a question, which kind of justification is relevant, and when to provide it. Some details of the current implementation are included.

We report here on a significant new set of capabilities that we have incorporated into our language generation system MUMBLE. Their impact will be to greatly simplify the work of any text planner that uses MUMBLE as its linguistics component since MUMBLE can now take on many of the planner's text organization and decision-making problems with markedly less hand-tailoring of algorithms in either component. Briefly these new capabilities are the following:

- (a) ATTACHMENT: A new processing stage within MUMBLE that allows us to readily implement the conventions that go into defining a text's intended prose style, e.g., whether the text should have complex sentences or simple ones, compounds or embeddings, reduced or full relative clauses, etc. Stylistic conventions are given as independently stated rules that can be changed according to the situation.
- (b) REALIZATION CLASSES are a mechanism for organizing both the transformational and lexical choices for linguistically realizing a conceptual object. The mechanism highlights the intentional criteria which control selection decisions. These criteria effectively constitute an "interlingua" between planner and linguistic component, describing the rhetorical uses to which a text choice can be put while allowing its linguistic details to be encapsulated.

The first part of our paper (sections 2 and 3) describes our general approach to generation; the rest illustrates the new capabilities through examples from the UMass COUNSELOR Project. This project is a large new effort to develop a natural language discourse system based on the HYPO system (Rissland & Ashley 1984), which acts as a legal advisor suggesting relevant dimensions and case references for arguing hypothetical legal cases in trade-secret law. At various relevant points we briefly contrast our work with that of Appelt, Danlois, Gabriel, Jacobs, Mann and Mattheissen, and McKeown and Derr.

Tense, Aspect and the Cognitive Representation of Time

Kenneth Man-kam Yip

Artificial Intelligence Laboratory
Massachusetts Institute of Technology

545 Technology Square
Cambridge, MA 02139

pp. 806-814

This paper explores the relationships between a computational theory of temporal representation (as developed by James Allen) and a formal linguistic theory of tense (as developed by Norbert Hornstein) and aspect. It aims to provide explicit answers to four fundamental questions: (1) what is the computational justification for the primitives of a linguistic theory; (2) what is the computational explanation of the formal grammatical constraints; (3) what are the processing constraints imposed on the learnability and markedness of these theoretical constructs; and (4) what are the constraints that a linguistic theory imposes on representations. We show that one can effectively exploit the interface between the language faculty and the cognitive facilities by using linguistic constraints to determine restrictions on the cognitive representations and vice versa.

Lexical Ambiguity as a Touchstone for Theories of Language Analysis

Lawrence Birnbaum

Yale University
Department of Computer Science
New Haven, CT 06520

pp. 815-820

This paper assesses several broad approaches to language analysis with respect to the problem of lexical ambiguity. The impact of the problem on both syntactic and semantic analysis is discussed, and several common methods for disambiguation, including the use of selectional restrictions and scriptal lexicons, are analyzed. Their shortcomings illustrate the need for complex inference to resolve ambiguity, which forms one of the key functional arguments in favor of integrating language analysis with memory and inference. However, it has proven surprisingly difficult to realize such an integrated approach in practice: An assessment of lexical disambiguation within some recent models which attempt to do so reveals that they rely largely on the traditional techniques of selectional restrictions and

VOX – An Extensible Natural Language Processor*Ammon Meyers*

Artificial Intelligence Project
 Computer Science Department
 University of California
 Irvine, California
pp. 821-825

Partial Constraints in Chinese Analysis*Yiming Yang, Shuji Doshita, Toyooki Nishida*

Department of Information Science
 Kyoto University
 Sakyo-ku, Kyoto 606, Japan
pp. 826-828

Grammatical Functions, Discourse Referents and Quantification*Uwe Reyle*

Department of Linguistics
 University of Stuttgart
 West Germany
pp. 829-831

Discourse Structure and the Proper Treatment of Interruptions*Barbara J. Grosz*

AI Center, SRI International
 Menlo Park, CA 94025
 & CSLI, Stanford University
 Stanford, CA 94305

Candace L. Sidner

BBN Laboratories
 10 Moulton Street
 Cambridge, MA 02238
pp. 832-839

Evaluating Importance: A Step Towards Text Summarization*Danilo Fum, Giovanni Guida, Carlo Tasso*

Istituto di Matematica, Informatica e
 Sisemistica
 Università di Udine
 Udine, Italy
pp. 840-844

Understanding Analogies in Editorials

scriptal lexicons, with all their drawbacks. The difficulty is shown to stem primarily from the theories of memory and inferential processing utilized. The implications for recent approaches to language analysis based on connectionist mechanisms are explored. Finally, the requirements imposed by lexical disambiguation on theories of memory and inferential processing are discussed.

VOX is a Natural Language Processor whose knowledge can be extended by interaction with a user.

VOX consists of a text analyzer and an extensibility system that share a knowledge base. The extensibility system lets the user add vocabulary, concepts, phrases, events, and scenarios to the knowledge base. The analyzer uses information obtained in this way to understand previously unhandled text.

The underlying knowledge representation of VOX, called Conceptual Grammar, has been developed to meet the severe requirements of extensibility. Conceptual Grammar uniformly represents syntactic and semantic information, and permits modular addition of knowledge.

In this paper we describe a method using semantic constraints to reduce the ambiguities and generate case structure from phrase structure in Chinese sentence analysis.

Semantic constraints written on semantic markers indicate the plausible case structure. Different sets of semantic markers are chosen according to the purpose. A priority evaluation scheme steers the analysis towards the most plausible structure first, without trying all possibilities.

A new algorithm is proposed which transforms f-structures into discourse representation structures (DRSs). Its primary features are that it works bottom up, that it is capable of translating f-structures without pre-imposing any arbitrary order on the attributes occurring in it, and that it handles indeterminacy of scoping by using sets of translations. The approach sheds light on how an efficient interaction of different components of a natural language processing model can be achieved.

This paper reports on the development of a computational theory of discourse. The theory is based on the thesis that discourse structure is a composite of three structures: the structure of the sequence of utterances, the structure of intentions conveyed, and the attentional state. The distinction among these components is essential to provide adequate explanations of such discourse phenomena as clue words, referring expressions and interruptions. We illustrate the use of the theory for four types of interruptions and discuss aspects of interruptions previously overlooked.

The paper deals with the problem of evaluating importance of descriptive texts and proposes a procedural, rule-based approach which is implemented in a prototype experimental system operating in the specific domain of text summarization. Importance evaluation is performed through a set of rules which are used to assign importance values to the different parts of a text and to resolve or explain conflicting evaluations. The system utilizes world knowledge on the subject domain contained in an encyclopedia and takes into account a goal assigned by the user for specifying the pragmatic aspects of the understanding activity. In the paper some examples of the system operation are presented by following the evaluation of a small sample text.

The widespread use of analogy in human communication underscores the

Stephanie E. August, Michael G. Dyer
 Artificial Intelligence Laboratory
 Computer Science Department, 3531 BH
 University of California
 Los Angeles, CA 90024
pp. 845-847

Integrating Text Planning and Production in Generation

Eduard H. Hovy
 Yale University
 Artificial Intelligence Project
 2158 Yale Station
 New Haven, CT 06520
pp. 848-851

Be Brief, Be to the Point, ... Be Seated or Relevant Responses in Man/Machine Conversation

Anne Vilnat, Gérard Sabah
 GR22, Paris VI
 4, Place Jussieu
 75230 Paris Cedex 05
pp. 852-854

SAPHIR+RESEDA, A New Approach to Intelligent Data Base Access

Bernard Euzenat, Barnard Normier, Antoine Ogonowski
 ERLI
 72 quai des Carrières
 94220 Charenton, France
Gian Piero Zarri
 Centre National de la Recherche Scientifique, Paris & TECSIEL
 via Barnaba Oriani 32
 00197 Roma, Italy
pp. 855-857

RESEARCHER: An Experimental Intelligent Information System

Michael Lebowitz
 Department of Computer Science
 Computer Science Building
 Columbia University
 New York, NY 10027
pp. 858-862

A Parallel-Process Model of On-Line Inference Processing

Kurt P. Eiselt
 Irvine Computational Intelligence Project
 Computer Science Department
 University of California
 Irvine, CA 92717
pp. 863-869

New Approaches to Parsing Conjunctions Using Prolog

need for a system which can recognize and understand analogies. This paper presents a theory of analogy recognition and comprehension, using as a domain letters to the editor of a weekly news magazine. Some of the issues facing a system which understands analogies in this domain are identified, initial work on this program is reviewed, and work in progress is discussed.

While the task of language generation seems to separate quite naturally into the two aspects of language generation (text planning and text production), it is necessary to have the planning and the production interact at generator decision points in such a way that the former need not contain explicit syntactic knowledge, and that the latter need not contain explicit goal-related information. This paper describes the decision points, the types of plans that are used in making the decisions, and a process that performs the task. These ideas are embodied in a program.

In the dialogue part of our system, we have tried to increase the user's possibilities to criticize the machine's results and require explanations of them. The system must then provide a clear justification: either by furnishing the chain of reasoning or by asking a "good question" when it failed. The system must also be able to engage in a real dialogue, with more than one question/one answer. To do that, the system must build and use several kinds of representation: the reasoning, the topics and a model of the user, which is used to tailor the system's responses.

This paper describes a transportable natural language interface to databases, augmented with a knowledge base and inference techniques. The inference mechanism, based on a classical expert system's type of approach, allows, when needed, to automatically convert an input query into another one which is "semantically close". According to RESEDA's theory, "semantically close" means that the answer to the transformed query implies what could have been the answer to the original question. The presented system integrates natural language processing, expert system and knowledge representation technology to provide a cooperative database access.

The development of very powerful intelligent information systems requires the use of many techniques best derived by studying human understanding methods. RESEARCHER is a system that reads, remembers, generalizes from, and answers questions about complex technical texts, patent abstracts in particular. In this paper we discuss three current areas of research involving RESEARCHER – the generalization of hierarchically structured representations; the use of long-term memory in text processing, specifically in resolving ambiguity; and the tailoring of answers to questions to the level of expertise of different users.

This paper presents a new model of on-line inference processes during text understanding. The model, called ATLAST, integrates inference processing at the lexical, syntactic, and pragmatic levels of understanding, and is consistent with the results of controlled psychological experiments. ATLAST interprets input text through the interaction of independent but communicating inference processes running in parallel. The focus of this paper is on the initial computer implementation of the ATLAST model, and some observations and issues which arise from that implementation.

Conjunctions are particularly difficult to parse in traditional, phrase-based grammars. This paper shows how a *different* representation, not based on

Sandiway Fong, Robert C. Berwick

Artificial Intelligence Laboratory
MIT, 545 Technology Square
Cambridge, MA 02139
pp. 870-876

tree structures, markedly improves the parsing problem for conjunctions. It modifies the union of phrase marker model proposed by Goodall (1984), where conjunction is considered as the linearization of a three-dimensional union of a non-tree based phrase marker representation. A PROLOG grammar for conjunctions using this new approach is given. It is far simpler and more transparent than a recent phrase-based extraposition parser for conjunctions by Dahl and McCord (1984). Unlike the Dahl and McCord or ATN SYSCONJ approach, no special trail machinery is needed for conjunction, beyond that required for analyzing simple sentences. While of comparable efficiency, the new approach unifies under a single analysis a host of related constructions: *respectively* sentences, right node raising, or gapping. Another advantage is that it is also completely reversible (without cuts), and therefore can be used to generate sentences.

On the Use of a Taxonomy of Time-Frequency Morphologies for Automatic Speech Recognition**Renato De Mori, Mathew Palakal**

Concordia University
Department of Computer Science
1455, de Maisonneuve Blvd.
Montreal, Quebec, Canada H3G 1M8
pp. 877-879

A computer vision approach based on skeletonization and hierarchical description of speech patterns is proposed. Learning hierarchical descriptions of phonetic events is discussed. Experimental results are reported showing the power of the approach in the recognition of diphthongs in connected letters and digits.

Reversible Automata and Induction of the English Auxiliary System**Robert C. Berwick, Samuel F. Pilato**

MIT Artificial Intelligence Laboratory
545 Technology Square
Cambridge, MA 02139
pp. 880-822

In this paper we apply some recent work of Angluin (1982) to the induction of the English auxiliary verb system. In general, the induction of finite automata is computationally intractable. However, Angluin shows that restricted finite automata, the *k-reversible* automata, can be learned by efficient (polynomial time) algorithms. We present an explicit computer model demonstrating that the English auxiliary verb system can in fact be learned as a 1-reversible automaton, and hence in a computationally feasible amount of time. The entire system can be acquired by looking at only half the possible auxiliary verb sequences, and the pattern of generalization seems compatible with what is known about human acquisition of auxiliaries. We conclude that certain linguistic subsystems may well be learnable by inductive inference methods of this kind, and suggest an extension to context-free languages.

DP-Matching: With or Without Phonemes?**Shigeyoshi Kitazawa**

Shizuoka University
Hamamatsu-shi, 432, Japan

Masa-aki Ishikawa, Shuji Doshita

Kyoto University
Sayo-ku, Kyoto-shi, 606, Japan
pp. 883-885

Attempts at automatic speech recognition have known several waves.

Early efforts were based on the faith that speech is a string of phonemes that can be isolated and recognized one by one. This wave broke when it became clear that the physical realization of a phoneme is smeared in time and mingled with that of its neighbors, and also context and speaker-dependent.

Next came the invention of the highly successful time-warping DP-matching methods, in which whole words are matched by templates. This wave is still going strong, at least in Japan, but it may have reached a high mark.

To probe this question, we investigate the case of the "jion", a subset of character readings that "generates" a large subset of Japanese. This set has low redundancy and contains many minimal pairs. Error analysis of DP-matching shows that most errors occur between pairs that differ only in their initial consonant, especially if it belongs to groups such as plosives or nasals.

Combining DP-matching with limited-scope phoneme recognition could break through present limits.