American Journal of Computational Linguistics Microfiche 64

the FINITE STRING

NEWSLETTER OF THE ASSOCIATION FOR COMPUTATIONAL LINGUISTICS

VOLUME 14 - NUMBER 2 MAY 1977

AMERICAN JOURNAL OF COMPUTATIONAL LINGUISTICS is published by the Association for Computational Linguistics.

SECRETARY-TREASURER: Donald E. Walker, Stanford Research Institute, Menlo Park, California 94025.

EDITOR: David G. Hays, 5048 Lake Shore Road, Hamburg, New York, 14075

EDITORIAL ASSISTANT: William Benzon

American Journal of Computational Linguistics *Microfiche 64 : 2*

REPORT ON 1977 ACL PANEL ON SPEECH UNDERSTANDING AND COMPUTATIONAL LINGUISTICS: A CRITICAL EXAMINATION OF THE A R P A PROJECT

STANLEY R. PETRICK

Jonathan Allen, Chairman
Jack Mostow, Carnegie-Mellon University
Donald Walker, Stanford Research Institute
William Woods, Bolt Beranek and Newman

The present report summarizes the statements of the panelists and paraphrases the questions and answers that followed their talks.

REPORT OF THE ACL 1977 ANNUAL MEETING PANEL ON SPEECH UNDERSTANDING AND COMPUTATIONAL LINGUISTICS A CRITICAL EXAMINATION OF THE ARPA PROJECT

S. R. Petrick

IBM T J. Watson Research Center

This panel discussion was chaired by Jon Allen of MIT. The participants were Jack Mostow of Carnegie-Mellon University (ably substituting for Raj Reddy on short notice), Don Walker of the Stanford Research Institute, and Bill Woods of Bolt, Beranek and Newman The panel format featured a half hour talk by each panelist followed by a period of discussion centered around questions from the audience.

Jon Allen began by reminding the audience that the ARPA five year speech understanding project had just ended last October, and that this was an appropriate time to step back and ask just what had been learned as a result of that project He then turned the podium over to the panelists to address that issue.

Don Walker began by supplying some general background information on all the ARPA-sponsored efforts before he turned to the specific accomplishments of SRI. They were all concerned with speech understanding rather than speech recognition, and they each focussed on one or more specific tasks SRI, for example, was concerned with information retrieval with respect to a navy ship data base. This concentration on specific tasks limited generality somewhat. There was also a focus on producing an operational system. This was good in some respects, but this emphasis on results did make it impossible to follow many promising research paths. Performance requirements, such as size of vocabulary and processing time, were not all realistic.

The approaches followed differed widely from group to group These diverse approaches were perhaps appropriate to the state of knowledge which existed, but they did limit the opportunity to share results. There were competitive aspects to the ARPA effort. In addition, the necessity to attend lots of steering committee meetings hindered local progress.

The computing power made available for the ARPA-sponsored efforts was good but still not enough It was not as much as the investigators had hoped for. Walker referred to the money which had been expended and remarked that although it was substantial, it was not enough. Not only would more computing power have been useful but more people and more time would also have contributed to improved results.

Turning then to a more detailed examination of the SRI effort, Walker discussed four system issues which had been dealt with: (a) integration of speech, syntax, semantics, etc., i e., specifying contributions and interactions of multiple knowledge sources perspicuously and efficiently, (b) cooperation, i.e., sharing information and avoiding duplication of effort, (c) evaluation, i.e., rating alternatives and choosing what to do next, and (d) attention, i e., avoiding dissipation of resources over competing alternatives He also discussed some component issues including those required for natural language specification, task modeling, and dialog interactions

Typical sentences in the SRI information retrieval application were cited and discussed These included

WHAT IS THE LENGTH OF THE LAFAYETTE?

THE GEORGE WASHINGTON?

DID GENERAL DYNAMICS BUILD THE SHARK?

WHO OWNS IT?

WHICH FRIGATES WERE BUILT BY CAMMELL LAIRD COMPANY?

The SRI system organization was discussed briefly Those portions of the system concerned with acoustics, phonetics and phonology were handled by SDC on their own computer while higher levels such as syntax and semantics were treated by SRI Basically, a flexible control facility provided for such alternatives as either building structure up from acoustically determined words or building it top-down in a syntax-directed fashion. In any case, a semantic network was constructed using scores of the degree of phonetic match involved with various alternatives.

Turning to the definition of the input language, Walker stated that it is accomplished by means of linguistically motivated (ATN) rules that are general and extensible over a variety of domains. These provide a means for adjusting ('tuning') the language definition to particular domains without loss of generality. Syntactic, semantic, and discourse information is combined within the rules that define words and phrases.

The SRI semantic component makes use of partitional semantic networks. It handles higher-order logical predicates, especially quantifiers It provides deduction routines for retrieval and inference that can access a supplementary relational data base in responding to a user's query It also provides a network substructive that is converted to an English sentence or phrase to answer a user's question.

SRI discourse modeling is based on in-depth studies of domain-oriented dialogs A model of dialog context is encoded through the use of semantic partitions. Dialog context (through semantic closeness and syntactic considerations) is also used to find the meanings of elliptical expressions and the referents of definite noun phrases

System integration was briefly discussed. In the SRI system this provides for interaction of information from various sources of knowledge -- syntax, semantics, and discourse -- as part of the language definition itself The SRI effort avoided commitment to a particular system control strategy, allowing instead for flexible use of various strategies for putting together words and phrases out of incomplete and uncertain fragments.

The SRI system control facility was the last topic discussed by Walker. The points made about it included the following: Ratings are assigned to incomplete sentence structures reflecting the highest possible continuation of each such structure. The priorities of available alternative continuations are assigned as a result of those ratings. Data structures are organized for testing hypotheses about utterances in a manner that avoids duplication of effort. Combinations of top-down, bottom-up, and bidirectional strategies are allowed. The flexibility provided by the SRI system control made possible extensive experimental studies to evaluate design alternatives. Control strategies which were evaluated included the following four paired alternatives:

1) Island driving versus left to right processing of utterances
2) Testing for all words versus testing for just good matches
3) Using sentential context versus ignoring it
4) Focussing on selected alternatives versus continuing on a single successful path until it succeeds or fails.

Experimental findings from these studies of alternatives included

a) Context checking is important; it helps priority setting
b) Testing for all words improves performance but it is costly. A more efficient mapper is needed.
c) Island driving produced wild exponential branching on long sentences.
d) Focussing on selected alternatives is less effective than best first following of a single path. Nevertheless, simple best-path-first parse algorithms are not good enough
e) There is a strong need for a phrase mapper for word coarticulation
f) A high false alarm rate was reported. Hit scores were better than false alarm scores

Bill Woods, the next speaker, began with a discussion of a speech spectrogram He described early experiments in which two investigators, Ken Stevens and Dennis Klatt, attempted to "read" spectrograms either on a segment by segment basis or making additional use of their knowledge of English phonology and syntax. They correctly transcribed between 70 per cent and 75 per cent of the individual segments considered in isolation (some segments were only partially identified, i.e., certain of their distinctive features were correctly identified), but attained 96 per cent word accuracy in an experiment which used a computer program as an aid to human classification together with the human skills and knowledge of linguistic constraints they commanded.

A number of observations were made. They included the following:

(1) Small function words are highly unreliable anchors.

(2) Acoustically based matching alone is insufficient; accidental word matches outnumber correct ones.

(3) Speech understanding is essentially a nondeterministic process.

(4) Sequential left-to-right scanning presents problems; the first and last words are often not pronounced carefully; the first stressed word is a more reliable anchor There is a need to recover from any garbled word, especially the first one

(5) The space of alternatives must be open-ended You might have to go way down the list of likely alternatives to find the right one

(6) There is a strong need for merged representations. Many similar hypotheses share common parts.

A discussion of the sources of knowledge in the BBN System followed.Acoustic-phonetic, phonological, lexical, syntactic, semantic, discourse, and factual sources of knowledge were considered.

A program to assign scores to each of a number of overlapping segments for each possible candidate phoneme was described. Results of an experiment were described in which the correct phoneme was identified 53 per cent of the time and the first or second choice was correct 70 per cent of the time. Per cent of phonemes correct within the first N choices for various values of N was reported as follows

N	1	2	3	4	5	6	7
% correct	53	70	75	79	83	84	86

The BBN verification component was next discussed It takes an analysis-by-synthesis approach to word matching. A synthesis-by-rule procedure is used to generate word templates, and these templates are compared with a parametric version of the signal to determine the closeness of match. The pragmatic ATN-based grammar which served as the syntactic component was identified as that described by G. Brown in a paper presented earlier that day.

A diagram of the BBN HWIM System organization was discussed briefly. The basic operation of the system was described as follows: (1) create initial theories (2) rank them, (3) give them to the parser, (4) have the parser suggest the n best words (5) verify best matches, (6) add results on to islands previously identified to form larger islands, and (7) let these new events trigger new theories and continue this cycle.

Woods then presented transparencies citing in parallel columns the ARPA-imposed 1971 goals and the degree of success attained by the BBN HWIM System in meeting those goals. I include reproductions of those transparencies.

	ARPA Goals (from 1971)	HWIM (1976)
	The system should:	The system did:
1	Accept continuous speech.	Accept connected speech
2	From many	From 3
3	Cooperative speakers of the general American dialect	Cooperative male general American speakers
4	In a quiet room	In a somewhat quiet computer terminal room
5	Over a good quality microphone	Over an ordinary close-talking talking microphone
6	Allowing slight tuning of the system per speaker	With no tuning of the system for the individual speaker
7	But requiring only natural adaptation by the user	And requiring no speaker adaptation
8	Permitting a slightly selected vocabulary of 1000 words	On 1097 words with no post-selection
9	With a highly artificial syntax	With a somewhat natural combined syntax and semantics (average branching > 100)
10	And a highly constrained task	And a well defined task.

11	With a simple psychological model of the user	With no user model
12	Providing graceful interaction	And modest interaction capabilities
13	Tolerating less than 10%	With 57% semantic error semantic error
14	In a few times real time	In about 1350 times real time
15	On a 100 Mips Machine	On a .35 Mips PDP-10
16		With 256k 36-bit words
17		With a hierarchical system organization
18		(Cost per utterance)
19	And be demonstrable in 1976 with a moderate chance of success	And was operational 12 October 1976

I also include information from a summary of results transparency which casts further light on the 57 per cent semantic error figure cited in item 13. In discussing this summary Woods stressed that the difficulty of recognition is very much affected by the branching factor associated with the syntactic component in question and with the distinctiveness of the paths allowed.

Summary of Results 10/31/76

124 Utterances by 3 Speakers

Average Utterance = 6.1 words = 1.8 seconds

	BIGDICT (1097 words)	MIDDICT (409 words)
Correct	41%	48%
Semantically OK	2%	4%
Incorrect		
Close	23%	20%
So-so	10%	6%
Bad	5%	6%
Gave up (150 theories)	18%	14%
System Broke	1%	1%
Estimated average branching (words)	196	67
Speed (times real time)	1350	1050

Woods closed with a discussion of those natural language problem areas which had been affected by results obtained through speech understanding research. They included the following:

(1) Coping with ambiguity (BBN did not eliminate this by fiat)

(2) Error recovery (island driving is good in this respect)

(3) Parsing algorithms

(4) Interaction of syntax, semantics, and pragmatics

(5) Degrees of "grammaticality" (BBN was just getting into this. SRI somewhat more so.)

(6) Grammar-writing methodologies (e.g., general categories and specific filters at SRI, domain-specific categories at BBN, and a compromise between these as suggested in the paper by Bobrow and Bates)

(7) Discourse and dialog models

(8) Human performance and linguistic theory.

The final talk of the afternoon was given by Tom Mostow of Carnegie-Mellon University. The CMU project featured a task domain concerned with information retrieval for a set of AI abstracts. Vocabulary size was 1011 words and a branching factor of 9.53 was cited, indicating a considerably more restricted syntax than that of the other efforts. Two distinct systems, HARPY and HEARSAY-II were implemented, and the results cited for each were:

HARPY

97% word accuracy

91% sentence accurac

95% semantic accuracy

27.9 Mips (about 30 times real speech rate)

HEARSAY-II

86% word accuracy

73% sentence accuracy

91% semantic accuracy

85.0 Mips (85 times real speech rate)

After a discussion of how signal variability and imperfect knowledge lead to a combinatorial explosion of alternative transcriptions, making speech understanding very difficult, Mostow gave more detailed descriptions of both HEARSAY-II and HARPY.

HEARSAY-II is an asynchronous, data-directed, hypothesize and test model in which cooperating, independent knowledge sources communicate by means of a "blackboard". In this model it is easy to deactivate a particular knowledge source, and this made it relatively easy to evaluate the worth of each such source. Focus of attention mechanisms were studied. The system is island-driven and makes use of a task-specific pattern matching grammar

HARPY is quite different. It integrates knowledge sources into a uniform state-space network at the phoneme level. All knowledge sources, high and low, must be integrated into such a network. It was claimed to be very efficient for sufficiently restricted languages with a sentence accuracy over 90 per cent on a 1 MIP PDP/10 in 10 to 30 times real time Proceeding to a new task requires that a similar but new grammar be produced, which Mostow claimed was not a major undertaking.

A general observation was made and illustrated by specific experience accrued at CMU. This was that different uses of the same linguistic knowledge require different knowledge representations. Precompiling knowledge into an appropriate data structure for a given use is more efficient than interpreting a single global knowledge structure.

The philosophy of language design adopted by CMU was, given a set of typical sentences, to construct a seqnence of increasingly larger languages each of which properly contains its successor. The last member of this sequence must account for the given corpus. In other words it was proposed that full linguistic constraints be relaxed at first to produce a simpler model with a correspondingly more efficient recognition procedure; it was argued that this provides a cheap filter on data, and the output of such a model can, in turn, be constrained by a more elaborate model. In illustration of this approach with respect to HARPY, it was suggested that a finite state net could serve as the simpler model and an ATN grammar as the more elaborate model. A HEARSAY-II example was also discussed in which a Markov model serves as the simpler model and a pragmatic template grammar as its more elaborate extension.

The problem of understanding ungrammatical speech was touched on briefly. The approach suggested was to use a template-based grammar but to accept approximate template realizations. If the normal grammar fails, find grammatical fragments and concatenate them if permitted by general rules.

Search techniques for both HEARSAY-II and HARPY were discussed. The locus search method of HARPY was cited as one of the main contributions of CMU research along with the HEARSAY-II blackboard model and the results obtained on the interpretation of ungrammatical utterances.

Future directions included extension of the system(s) to provide for more general (larger and more flexible) input languages and extension of the methods developed to such other AI tasks as the development of vision systems.

Following the prepared panel talks questions were directed to the panel members from the floor. I cannot identify all the questioners so I will not attempt to supply names. I have paraphrased questions and answers in most cases.

Q: What linguistic constraints were imposed other than a 1000 word vocabulary?

A: Formal grammars of one type or another were used (for example, an ATN grammar for the BBN effort). Sentences which were grammatical with respect to these formal grammars were read by subjects.

Q: Was any use made of intonation?

A: (Woods) Some use of intonation constraints was implemented into the HWIM System rather late in the program but was never tested in the final crunch.

Q: What results from speech work should help in text processing?

A: (Woods) The development of a model for completing unfinished sentences. This has certain psycholinguistic implications.

A: (Walker) Language definition involving general categories and pragmatic grammars for one thing. The integration of all knowledge sources at the phrase level is also a significant contribution.

A: (Mostow) The determination of how much milage can be squeezed out of cheap methods. We expect to continue this type of speech research but not to do much more on high level constraints.

Q: Why did CMU do so much better than the others?

A: (Combined responses of all three panelists)

CMU picked a simpler domain perhaps and definitely used a much cruder (more restrictive) grammar.

The BBN LISP System and operating system were less efficient than CMU's lower level programming.

BBN aimed for a tougher target (as evidenced by their more realistic grammar and correspondingly higher branching factor) and didn't make it in spite of a ten % per month improvement over the last few months.

The BBN results were intermediate bench marks along the way toward a target which was never reached because of a lack of time whereas CMU was more realistic in budgeting its time.

CMU had no real theory. It just threw in a new dictionary entry whenever necessary.

CMU emphasis was placed on general AI techniques, not linguistic principles

BBN could only process 10 utterances per night whereas CMU could process 10 utterances per hour. Hence, CMU could consider more data SRI did not have a total system avialable to them, being dependent on SDC for their front end, so they were at a disadvantage in this respect. All speakers agreed that once real time speeds are achieved, the user interaction possibilities will be attractive.

SRI not only had no opportunities for fine tuning, but their approach was claimed to be even more linguistically motivated than BBN's. It was similar to BBN's before the latter's switch to a pragmatic grammar.

SRI reported a 40/1 false alarm-to-hit ratio on the part of the acoustic component supplied by SDC but still achieved close to 90% semantic success Its system was estimated to run in 200-300 times real time on a vocabulary of 320 words and a grammar with a static branching factor of 110.

Precomputation gives HARPY its efficiency but incorporation of new knowledge into HARPY makes new precomputation necessary and it is expensive.

Q: Does anyone currently have a government contract to pursue speech recognition?

A: Not to anyone's knowledge, but more general contracts may include some speech work. Speech understanding is not dead but is currently in abeyance.

Q: Is the rumor true that NSF may pick up some speech work and therefore some projects that would otherwise be funded will suffer?

A: It is true in the sense that people previously supported by ARPA funds are entering the competitive pool for a static amount of research funds There is a "new player" however in the form of the Sloan Foundation, which is supporting less capital intensive research, is less systems-oriented, and is seeking to encourage cooperation between different groups.

Q: What would you do different if you were starting over?

A: (Woods) I would do some things different but probably that would be wrong. I would give more reassurance to new people starting work on the project that they were on the right track. ARPA managers need such reassurance too. More intermediate results might have kept up the interest of the sponsors. Too many decisions had to be made on the basis of inadequate experimentaion.

A: (Walker) Groups were brought together and educated and then disbanded.

Q: I can hear the difference between /s/ and /z/, etc. Why don't you do some better acoustic work so as to be able to make such discriminations reliably?

A: You "hear" the difference in large part by means of higher constraints such as syntax, semantics, and even pragmatics. It is well documented that the information for identification just isn't in the segment to be recognized in many cases.

Q: Are there any companies now working on speech recognition?

A: Yes. IBM is, and other companies such as Threshold Technology are working on word template matching.

American Journal of Computational Linguistics *Microfiche 64 : 16*

ESSAYS ON LEXICAL SEMANTICS

VOLUME II

EDITED BY V. JU. ROZENCVEJG

SPRÅKFÖRLAGET SKRIPTOR
Box
S 104 65 Stockholm
1974

282 pages
SKr 60

ISBN 91-7282-065-9

REVIEWED BY ERNST VON GLASERSFELD
Department of Psychology
University of Georgia
Athens 30602

Since the editor did not provide an introduction to this second volume, the only information we get about the authors of the 9 papers it contains is in a short blurb on the back cover of the book.

> The young linguists who gathered fifteen years around the Laboratory of Machine Translation [in Moscow] were attracted by the prospects of creative endeavor which the problem of automatic transaltion offered: the problem itself and the search for a way to solve it did not permit any dogma to dominate... They wanted...to understand and clearly formulate how a person understands the meaning of an utterance in the given language and then constructs a message in another language that preserves the original meaning.

Reading the papers, it quickly becomes clear that, though these authors may indeed not have a common homogeneous theory of semantic analysis, they nevertheless are what I should call "dogmatic" in their approach to meaning and how it functions in communication. Most of them refrain from making their basic philosophical position explicit; but, as I read them, I get the impression that they would all subscribe to the programmatic statement on the first page of the first essay in the volume, Towards a Functioning 'Meaning-Text' Model of Language, by I. A. Mel'čuk and A. K. Žolkovskij:

> The main aspect of human language behaviour consists undoubtedly in the speaker's transmitting to the hearer certain information, or content, which the present authors decided to refer to (for brevity's sake and rather indiscriminately) as MEANING. The hackneyed phrase "Language is a means of human communication" implies nothing beyond the fact that the speaker has a certain meaning in mind and produces a text to convey it, and the hearer extracts the intended meaning from the text. Language (de Saussure's langue) functions here literally as a mechanism, a 'meaning-text-meaning' transformer. The linguistic model should therefore aim at 'meaning-text' transition as its core operation.

At the risk of appearing pedantic, I am going to look rather closely at this passage. It purports to lay out the foundation on which models of language can be built and to establish the point of view from which we are to examine language. It is easy to agree that "language is a means of human communication." But if we agree to that - and I, for one, certainly do - we shall be obliged to take into account what we know about communication and its limitations. And that, I would suggest, entails a certain amount of disagreement with what is said in the rest of the quoted passage.

From a communication-theoretical point of view, to say that text conveys meaning is at best a metaphor; and to say that the receiver "extracts the meaning from the text" is a misrepresentation. Even in simple telegraphic communication the receiver does not, indeed could not, extract the meaning of the dot-and-dash signals from those signals. The fact that, say, a single dot means "e" can be extracted only from the conventionally established listing that links the set of signals to the letters of the alphabet. Such signals have been correctly described as "selectional instructions"; that is to say, they tell the receiver which letters he is to select from the list of letters which he already has. A signal that has no conventional link (semantic nexus) to a letter in the code list simply has no "meaning". Although on the level of telegraphy signal meaning is clearly a much simpler affair than the meaning of sentences and texts (where all sorts of combinatorial rules come into play), one important aspect remains the same for any form and any level of communication: Meaning is under all circumstances something that the receiver has to construct or, if you will, reconstruct, out of material he already possesses. The signal, be it an electric impulse or a string of phonemes, may tell him what to choose and how to put it together, but it can never supply the material itself. To put it very simply: Meaning does not travel between a source and a destination or between a sender and a receiver; what travels is a signal - and if communication is to be achieved, that signal must be semantically linked to the same meaning at both ends. This immediately raises the question: How can we ever be sure that the sender's meaning and the receiver's meaning are, in fact, the same? To answer that question we should need more than a theory of language we should need a theory of knowledge, i.e. a fully worked out epistemology.

Judging by the essays collected in this volume, none of the authors seems to be wholly aware of this indispensable connection between a theory of meaning and a theory of knowledge. This fuzziness leads to a rather indiscriminate mixing of acute practical observations and more or less inappropriate theoretical metaphors; and this, in turn, makes for difficult reading - if by "reading" we mean the endeavor to reconstruct what the author intended to say.

Mel'čuk and Žolkovskij come closest to my understanding of meaning when they say "meaning appears as a construct, a bundle of correspondences between actual (content-) equivalent utterrances..." or "the command of linguistic meaning manifests itself in the speaker's ability to express one and the same idea in a number of different ways ..." (p.2). The ability to paraphrase thus becomes a criterion for semantic competence and "designing functioning paraphrasing systems" becomes a major concern of semanticists. This, I believe, is indeed one way of successfully avoiding epistemological problems. Without trying to achieve an act of communication, the semantic analyst (provided he is a proficient speaker of the language) can be himself decide whether or not two or more arrangements of linguistic items (i.e., signals) are semantically linked to one and the same conceptual structure (idea). He can decide this by himself because the idea, the representation of the linguistic items, and the linkage between these linguistic and conceptual parts are all in his own head.

According to Mel'čuk and Žolkovskij "the first and principal goal of present-day semantics [is] the construction of a meaning-text model of language." (p.3) One might think that such a model would require a rather complete description of the "ideas" (i.e. the non-linguistic conceptual structures) on the one hand, and on the other, of the linguistic items (morphemes, words, syntactic structures) by means of which they can be "expressed". But the authors are careful to stress that they have singled out an area of semantics "which is linguistic <u>par excellence</u> and aims at modelling such a handling of meaning as involves nothing but a command of language and requires no special information about the surrounding world." (p.4) They add in parenthesis: Obviously, this distinction is hard to draw in many specific cases." Quite so. The distinction is indeed hard to draw because the two kinds of information are not different in the way they are purported to be different. "Command of language", after all, is knowledge of regularities and rules the language user gathers in his <u>experience</u> of language, and "information about the surrounding world" is knowledge of regularities and rules he gathers in other kinds of experience. According to one's epistemological position, the "surrounding world" will be involved in both or in neither.

The few semantic representations Mel'čuk and Žolkovskij illustrate (e.g. a relational pattern for the situation designated by expressions such as "to buy", "to sell", "to pay", etc.) would seem to be compatible with the rather more complete conceptual analyses that have led Schank (1975) and Fillmore (1975) to speak of <u>scripts</u> or <u>frames</u> respectively. But scripts and frames are conceptual scenarios derived from and dependent on actual experience, which is to say, they are representations of a surrounding world". Mel'čuk and Žolkovskij suggest that experiential knowledge ("a huge encyclopedia and 'axioms of reality' as well", p.17) is something essentially different that should be kept "in the background, but immediately available when necessary for all kinds of checking operations, etc." (p.17)

Their paraphrasing system, thus, is based not so much on underlying conceptual structures but rather on the exchange of <u>linguistic</u> items, an exchange which is to be carried out with the help of a large number of substitution lists that take into account certain relational properties of the lexical items they contain. These relational properties are referred to as the "syntax of semantic components" (p.44). Though the relations comprised in this "syntax" are many more than, say, the cases in a case grammar, they still seem to be derived from the structure of language rather than from the structure of non-linguistic conceptualization; that is to say, they reflect relations between items in language and only very indirectly, if at all, the conceptual relations by means of which we construct such entities as objects, activities, processes, and states.

In short, the "meaning-text model", because it approaches meaning "as the invariant of synonymic transformations" (p.41), will presumably lead to an extremely particularized mapping of what can be exchanged for what, on the level of linguistic expressions. In the last part of their paper the authors say that it "constitutes a step towards experimental linguistics" (p. 42) and that puts the effort into the proper perspective. Establishing what is and what is not "synonymous" in a given language and according to a given criterion of synonymy, is undoubtedly a valuable enterprise. But the discovery and mapping of synonymous expressions, acceptable paraphrases, and semantic equivalences, though it certainly requires linguistic competence, does not automatically provide a model for linguistic competence, let alone for linguistic communication. Hence I find it difficult to agree that this kind of investigation should be considered "the principal goal of present-day semantics.

In the second paper of the collection, Materials for a Russian-Somali Dictionary by A. K. Žolkovskij, the difficulty for readers who don't know Russian will be compounded if they are also ignorant of Somali. Though the translation, throughout the volume, seems to be quite adequate, the very fact that semantic analysis (inevitably) often leads to language-specific results, means that the reader has to take much of what is said on faith. Nevertheless even if a particular analysis is not applicable to the language or languages one knows, it may well stimulate ideas as to why and how certain items should be analyzed.

In this paper Žolkovskij "proposes tentative entries for a Russian-Somali dictionary consisting of 105 nouns occurring frequently in political texts and compiled according to the recommendations set forth above." (p.56) The dictionary entries take up 22 pages, i.e., two thirds of the paper. The brief "recommendations" contain an even briefer explanation of a particular translation problem: "Somali lacks the nominal transform of the verbal phrase." Hence, expressions such as _the policy of the colonialists_ have to be transformed into something like _the policy pursued by the colonialists_, before they can be translated into Somali. This is apparently complicated by the fact that Somali has a variety of "empty verbs" and rather strict rules of usage that determine with what items each of these can be used. The problem seems to be similar to the one English-speakers encounter when they want to take a bath in Italy, because in Italian one does not _take_ a bath but _makes_ it (_fare il bagno_). In Somali, however, the problem is compounded by the fact that the "empty" verbs mostly require many more specifications, such as to whom the object belongs, where it is, for what reason it is involved, etc. A semanticist would be interested to know whether there are any _conceptual_ regularities or rules to be gleaned from the differential usage of these "empty" verbs in Somali and whether, in fact, they are quite as empty as they seem. The paper does not attempt any such analysis but gives, for each Russian head word, a selection of Somali constructions that have to be used when the word is translated. Hence, this mini-dictionary will no doubt be useful to a student of Somali but it adds little to our knowledge of the semantics of Somali or any other language.

L. N. Iordanskaja makes a courageous attempt to analyze A Group of Russian Words Denoting Emotions. She briefly refers to the basic classifications of Spinoza (joy, sorrow, and desire) and Descartes (who added love, hatred, and admiration) and then makes a radical division between "emotional states" and "emotional attitudes". Though the criterion for this division is not explained, the definitions that follow suggest that "emotional states" are always considered the result of a specific stimulus-situation (actual or envisioned) while "attitudes" are not. Intuitively this seems a sensible distinction to make; but, I believe, it is somewhat misleading to proceed as though emotional states could be treated without any reference to attitudes. Indeed, one of the three characteristics the author uses to describe the meaning of the chosen Russian emotion-words is the polar dimension of "positive (pleasant)/negative(unpleasant)", and I would argue that in most, if not in all, cases it is precisely an "attitude" that makes an individual evaluate something as "pleasant" or "unpleasant". Apart from the limitation of the words that are being analysed, there is nothing to disagree with in the paper. As the author herself says, "many (of the definitions) may probably be refined." Anyone familiar with Robert Plutchik's (1962) classification of emotions will not find much novelty here.

In a 25-page essay, N. N. Leont'eva deals with the almost universal phenomenon of The Semantic Incompleteness of Texts, i.e. with the remarkable fact that human users of language are so very often able to construct a satisfactory meaning for linguistic expressions when these expressions supply only incomplete directives for such a construction. She differentiates syntactic incompleteness (ellipsis), the referential indeterminacy of pronouns (anaphoric expressions), the need for information that was given earlier in the text, lack of one of the elements required by the relational configuration of a particular concept (e.g. the knife cuts, where the agent is missing). and all the experientially established relationships, causal and otherwise, that texts presuppose in the reader.

Leont'eva has compiled a set of 41 "semantic relations" and half a dozen "metarelations" (relations relating relations) by means of which the semantic connectivity of a text can be mapped and tested for completeness. Each of the relations has two terms. Wherever the text indicates a relation but not both the terms it purports to relate, there is a case of semantic incompleteness and an occasion to ask how the language user manages to infer or retrieve the item that is lacking in the linguistic expression. This is clearly a very sensible pragmatic way to proceed because it leads on the one side to progressive refinement of the master-list of relations and, on the other, to insights into the as yet largely uncharted inferential operations the ordinary language user carries out when he "understands" a piece of language.

V. Ju. Rozencvejg in Notes on the Meaning of some French Words, does much the same as Fillmore (1965) in his analysis of "entailment rules". The main part of the paper deals with aller and venir while a short final section treats a wholly different area, namely French particles of affirmation and negation.

Unlike Fillmore, who did not include the combinations of go and come with the infinitive of another verb, Rozencvejg does consider the somewhat specialized infinitive constructions (e.g. aller voir, venir se placer) and, I believe, it is the analysis if these expressions that leads him to overrate the importance of the "directional" element. Thus he states that the sentence Olga va (de Moscou) à Grenoble "communicates" to the addressee that the speaker "is located at A (Moscow) or at some point between A and B (Grenoble), but not at B or a point behind B or beyond B ..." If this were the case, the language user would often have to have a much more detailed map of the world in his head than he actually has. In fact, a little later Rozencveig himself gives London as example of a point "behind" Grenoble as seen from Moscow; actually London and Grenoble are almost exactly equidistant from Moscow but more than a third of that distance apart from each other.

Later in his paper, Rozencvejg documents the important observation that it is precisely by the use of the spatial implications of these verbs of motion that a skillful writer influences the visual representations of his readers - "not unlike a cameraman who controls the perception of the viewer by appropriately positioning his camera." (p.154) There is no doubt that this is so, and it should alert the semantic analyst to the fact that the use of aller and venir (as in the case of go and come in English) is constrained rather by the addresse's location and direction of view than by the speaker's.

Semantic Analysis of Russian Negative Sentences, by J. V. Padučeva, illustrates in a semi-formal way the many difficulties encountered by the transformationalist in his attempt to formulate rules by means of which the various forms of negation in Russian (and some other languages) might be derived. Some 30 pages of classified examples (most of them in Russian, but some also in French) will provide a valuable survey for investigators of negation in other languages, especially if they can read Russian. The author ends by saying:

> "It is difficult to say whether one should, on the basis of the complex problems connected with this transformation, draw a pessimistic conclusion about the possibilities of a formal description of language, or should be satisfied with the depth of analysis required to solve even the most simple concrete problem within the framework of this theory." (p.202)

This reader would suggest a third way out, namely that the failure of transformational grammar may not necessarily imply that language cannot be formally described by other theories.

The last three essays of the volume, covering some 170 pages, are G. V. Dorofeev and Ju.S. Martem'janov and examine different aspects of the formalization of meaning and logical inferences within and between sentences. The 4-line abstract that precedes the first of the three papers tells us:

> The present paper follows up the series of semantic descriptions of groups of words pertaining to certain situations of reality. The notions described here are "obstacle", "possible", "way", "actual way", "leads to" and some others. (p.207)

This at once tempts me to ask: Well now, what will it be - descriptions of words in terms of notions or descriptions of notions in terms of words? - The next sentence, the first of the essay proper, asserts that "to describe an informal meaning formally means to substitute it completely, in all cases of its use, with a new, formalized concept." (p.207 Now I am really baffled. If we apply this definition of "to describe formally" to the above Abstract, we must conclude that in this paper the notions "obstacle", "possible", etc., are going to be substituted by "new, formalized concepts." Is that really what the authors intend to say? - As it turns out, it is not. What they are actually doing is, in a sense, the opposite and therefore quite sensible. They try to grasp the notion or concept of, say, "obstacle" and then to find a formalized (logical) expression that renders it without too much distortion. That is to say, they try to substitute the word by a formal expression in such a way that the notion or concept remains the same. For instance:

> The statement "Some phenomenon b is an obstacle to another phenomenon g" is equivalent to: "If event b occurs, event g does not take place", or, in the notation of mathematical logic:
>
> $$\underline{b} \rightarrow \neg \underline{g} \text{ .} \qquad \text{(p.209)}$$

Though this is by no means their final formal representation of "obstacle", it does show what is being substituted for what and that the final goal of the formalization is to be able to write some sequence of conventional logical symbols (such as the one given) for every occurrence of the word "obstacle" without doing too much violence to the concept or conceptual structure that constitutes its meaning. "Semantic description", in this paper then means the replacing of ordinary-language words by logical formulas which, intuitively, are deemed to be semantically equivalent, i.e., deemed to have the same conceptual content.

The authors note that an atemporal (classical) logic does not furnish expressions that could successfully be substituted for the words they are interested in. They introduce a time factor, and that allows them to come up with a rather flexible, comprehensive notation. But even in its final and most elaborate form, which involves set-theoretical expressions such as "closure" and "neighborhood", the formalization is still based on certain assumptions that seem unacceptable to a conceptual semanticist. To give one example, the concepts of "state" and "event" are throughout differentiated by the assumption that the first takes up a "period" (i.e., a stretch of time) while the second corresponds to a "moment" (i.e., a single point in time). It is easy to show that what we call an "event" involves under all circumstances a change, and that, conceptually, a change requires a sequence of at least two points. In fact, the difference between state and event is not the respective presence or absence of a lapse of time but rather that, with regard to a particular aspect, an item is considered the same or not the same at two points in time.

The words analyzed in this essay are, in the translation, listed in English only. That is unfortunate because the English word "obstacle", for instance, which serves as the main example, does not really fit the authors' formalization. In our language it would not be correct to equate the meaning of "obstacle" with that of "preclusion". If *b* is an obstacle to *g*, it would often be quite incorrect to formalize this state of affairs as "if *b*, then not-g". Obstacles need not preclude achievement, they merely make it more difficult. After all, we have competitions called "obstacle races" and they would not be much fun if, by definition, all competitors were precluded from reaching the goal.

In their second paper, using parts of the notation introduced in the first, the authors set out "to reconstruct the implicative and causal relations existing between the sentences of a text ..." (p.225) The text they use to demonstrate their method of analysis is a Russian fable about a fox and a crane. The highlights of the story are, (a) the fox invites the crane to a dinner of porridge which the crane cannot eat because it is spread on a plate, and (b) the crane then returns the invitation and serves soup in a pitcher into which the fox cannot squeeze his head. - It is clear that "obstacles" and their deliberate creation are of paramount importance in this tale. But the analysis goes very much farther. With the help of a large number of pre-established formalized representations (some two dozen of which are listed) for expressions that include difficult concepts such as "A is vindictive", "A is perspicacious", "A is a friend of B", and the rather startling "A is a 'feed-back system'", some 50 inferences are derived from 10 sentences of the story. These inferences, we are told, are what makes the text a connected whole I have no doubt that is what they do, but the explanations as to how the formalized expressions have helped in bringing out the connective implications are either inept or incomplete (or both). One of the reasons for this failure is that the reader finds it almost impossible to patch together the original unadulterated fabel. There are lists and tables with comments, coded expressions, translations, and the often wholly mysterious result of the analysis. (e.g. one "transition to deeper semantic structure in terms of the basic notions used", in the case of the sentence "The Crane ate up all the soup himself" is the opaque statement: "Fox is 'feed-back system'".)

It is clear that an enormous amount of work has gone into this analysis and this reader feels that much of it would be very interesting and useful if one could really get hold of it. It is a great pity that the two authors, who refer to Fillmore's Entailment Rules in a Semantic Theory, from which they certainly have gained inspiration, did not also gain some of Fillmore's exemplary clarity of exposition.

The third essay of this team of authors, Description of the Meaning of Words for the Purpose of Inference, makes a few distinctions that, I think, demonstrate the value of this kind of work. There is, among others, an illustration of the point that the linguistic negation of a statement cannot always be replaced by a simply logical negation; for if one says, for instance, "A did not succeed in q", there is, beyond the negation of success, the definite implication that A attempted g, and this implication is crucial for the proper comprehension or semantic analysis of the statement.

This last example may serve to make explicit a point which these three essays tend to obscure: It is not the formalization that leads to such semantic or conceptual insights but the analysis that is a prerequisite to formalization. Once analysis and formalization have been achieved, a suitable notation may help us to keep track of details of the analysis that might otherwise get lost again; but neither formalization nor notation can tell us more about the meaning of words than what a native speaker's semantic or conceptual analysis has brought out. What kind of formalization and notation we want to use will always depend on what we want to do with it. The strictly linear notation used here may be very suitable for certain purposes - for the human reader, however, a two-dimensional, more diagrammatic notation, such as for instance the one created by Roger Schank (1975), makes comprehension a good deal easier.

In conclusion, there is a lot of interesting material and much suggestive exploration in this book but I would recommend it only to readers who have the time and the patience to go through every passage more than once. Some intricacies of semantic analysis are difficult to follow no matter how lucidly they are presented. Here the difficulties for the reader are greatly increased by a general lack of clarity and organization in the exposition. Much of this reads like research reports prepared for constant collaborators or colleagues who are well acquainted with the work. Finally, since this book was reproduced directly from a type script, it is a pity that the typewriter used for the corrections was not the same as the original one and that so many errors managed to survive.

References

Fillmore, Charles J. Entailment rules in a semantic theory. Report No. 10. Project on Linguistic Analysis. Columbus, Ohio: Ohio State University, 1965.

Frame semantics and the nature of language. Annals of the New York Academy of Sciences (1976) 280, 20-32.

Plutchik, Robert The emotions: Facts, theories, and a new model. New York: Random House, 1962.

Schank, Roger C. SAM - A story understander. Research Report No. 43. New Haven, Connecticut: Yale Universitv. 1975.

CONSTITUENT AND PATTERN IN POETRY

ARCHIBALD A. HILL
Professor Emeritus of English and Linguistics
University of Texas, Austin

THE UNIVERSITY OF TEXAS PRESS
Austin 78712

xiv + 157 pages 1976 LC 75-32582
$11.95 ISBN 0-292-72010-6

REVIEWED BY JAMES JOYCE
Computer Sciences Division
Department of Electrical Engineering and Computing Science
University of California, Berkeley 94720

Constituent and Pattern in Poetry is a collection of Professor Archibald Hill's essays on the application of linguistic principles to short poems in English. The essays are revisions, sometimes minor and sometimes major, of previously-published studies. They are, Hill tells us, "concerned with poems as designs, and all fit into a thought-out scheme of the relations of literature and language." The view of language he brings to this collection is admittedly Bloomfieldian, and it might be argued that no amount of revision could bring his structural linguistics-based observations up to date. Yet his particular brand of linguistics is not a point of contention here; indeed, the book's strengths and weaknesses stem from other considerations.

The strength of the dozen essays making up the book is that an intelligent, sensitive reader of poetry is sharing insights with us. His sole dogma, it appears, is the invocation of the Joos Law, "to choose the least meaning for any unknown item and the maximal meaning for its context." The way this law is used reminds me very much of the basic principle of New Criticism, as well as what seems to me to be a basic tenet of scientific inquiry: that is, to make minimum assumptions about the data that yield maximal insights or organization. Usually the Joos law of semantics is applied with good sense and profit to us all. A second principle guiding Hill's reading of poetry is one I would like to see writ large over each departmental office and lounge in every university:

In a work we admire
we do not have to justify every line
to yet think it excellent.

The same might be said of linguistic studies without violating the spirit of the statement or its applicability.

There are two things in this collection that disturb me; one is what I feel is a design flaw, and the other is an unfortunate choice of a measure. In several of the essays diagrams are described or suggested in varying degrees of detail (p. 21, for example), but only one is produced. To be sure, they are not promised by the usual means of a reference to Figure 4, but their presence might have added to the reader's comprehension of material and perhaps would have saved Professor Hill the work of such detailed prose where a diagram might be better.

A second, more substantive point, concerns his method for deciding whether an image is invoked in a passage. Although admitting he is using "a rough kind of calculation" he appears not to realize just how inappropriate the measure is. In a discussion of Frost's poem "Bereft" he cites

> Leaves got up in a coil and hissed
> Blindly struck at my knee and missed.

and observes that a lexical item, such as hissed, either refers to a snake or does not. "The chance of a snake reference by pure chance is then one out of two. I am of course perfectly aware that an accurate weighting would be very difficult, probably impossible." (p. 96) The first assumption he makes is that whether the word hissed refers to a snake or not is a matter similar to whether a tossed coin comes up heads or not. That is, the probability that hissed refers to snake is 1/2, just as the probability that a tossed coin coming up heads is 1/2. Hill admits there is another way of calculating such probabilities, but continues with the one described above because it is easier. He then discusses the probability that several "snake items" occur in the same passage, taking the probabilities (each assigned 1/2) and multiplying them together to get (for four items) a probability of 1/16 that the references to a snake are by chance. Since Hill says the order of the "snake items" is significant, and it takes three items (out of the four) to make a sequence, the 1/16 figure should be multiplied by 1/6, giving a final answer of 1/96 that the snake references in the passage by Frost are by chance. I'm afraid Professor Hill's adviser on statistical matters has misperceived the phenomenon being modelled probabilistically and combinatorically, and Hill apparently has not understood fully what assumptions about language such a model would require to be applicable.

Also, regarding the order of the "snake items", Hill finds it significant that it is: "got up in a coil", "hissed", "struck", and "missed". For the order to figure into the significance of the four "snake items", there has to be a likelihood that the order of actions is natural among snakes. I discussed the matter with a herpetologist at San Francisco's Steinhart Aquarium, and learned that 1) most snakes do not hiss, 2) and those that do use hissing as a bluff rather than as a prelude to the attack. The herpetologist suggested Frost may have used hissing to substitute for, say, a rattlesnake's rattling. As a summary of his observations he suggested the passage from "Bereft" was "a matter of what you want to interpret it being." That is, to a herpetologist the passage is not a clearly accurate metaphor for snake behavior. Certainly snakes coil, hiss, strike, and sometimes miss their strike, but the particular order Hill cites as significant is an order that someone who knows snakes finds open to question.

Before the patient reader of this review turns off the microfiche reader in disgust at what appears a compounding of error by absurdity, let me state there is no question in my mind the passage at issue evokes a snake; I am sure there is no doubt in Hill's mind the passage evokes a snake. The difference between the two of us lies in the soundness or lack of soundness of the argument Hill uses to illustrate his point. The significance of the order of "snake items," *if taken into consideration as they reflect the activity of* snakes, does not count toward the probability that the image is that of a snake, but against it.

What Hill might have done instead is a cluster analysis of the "snake items" to see that the density of snake items is greater than the density of other kinds of images. One example of such use of cluster analysis is John Smith's work on imagery in Joyce's A Portrait of the Artist as a Young

Man, described in "Thematic Structure and Complexity," Style 9 (Winter, 1975), pp. 32-45. Neither the passage from "Bereft" nor the entire poem is large enough to support a statistical analysis, but that does not prevent us from applying the technique as a way of summarizing the data for us so long as we make no claims of probability. Numbers, after all, can only guide interpretation. That Hill's 1959 essay should know about Smith's 1975 essay is, of course, impossible; yet Hill's book is not just a reprint of the essays, but a revision as well. And cluster analysis has been known as a technique longer than either essay has been in print.

Yet Constituent and Pattern in Poetry is a book I found enlightening on a number of points, and I believe it is a valuable collection because of those insights. I very much liked Chapter 8, "The Locus of the Literary Work," and feel it should be better known in literary circles than it is. The chapter views a poem not as a work on paper, as both New Critics and Structuralists would have us believe, but as a pattern, a poetic structure analogous to langue; the interpretations of the poem, however, are not parole, but conclusions about parole (p. 86). It is the text (or various texts of the poem, which need not be written to be texts) that Hill identifies as corresponding to parole in this case. A diagram of this might be:

```
 ---------------
 | The poem, or |
 |   Langue     |
 ---------------
        |
 ---------------
 | The texts, or |
 |    Parole     |
 ---------------
        |
 ------------------
 | Interpretations,|
 | or Conclusions  |
 | about Parole    |
 ------------------
```

I would want to add another part to this picture, one I hope Professor Hill would not object to:

```
 ---------------
 | The poem, or |
 |   Langue     |
 ---------------
        |
 --------------------
 | The poet's oeuvre,|
 |  as context for   |
 |    the poem       |
 --------------------
        |
 ---------------
 | The texts, or |
 |    Parole     |
 ---------------
        |
 ------------------
 | Interpretations,|
 | or Conclusions  |
 | about Parole    |
 ------------------
```

The poet's *oeuvre* exists between the poem as pattern in the poet's mind and the texts of a poem's realization. The *oeuvre* provides a context for composition at the moments of creation; how many moments those number, and whether they are contiguous or not, is irrelevant here. The various texts of a given poem, whether versions of one person's poem or from an oral tradition, begin as patterns in the mind and find realization as text influenced in composition by the poet's _oeuvre_. Critics make conclusions about _oeuvre_ when they consult a concordance for data about a word or do stylometric studies. Such studies ultimately allow us to interpret a particular text by providing accurate models of the poet's _oeuvre_; that is, we form better conclusions about _parole_. Such studies as well allow us to make inferences about _langue_ if that is the target of our study.

Hill states as one of the general ideas he holds strongly that "the method of literary investigation, when properly carried out, must not only rest on a linguistic base but also go beyond it." (p. x) This seems to me to be so reasonable and appealing I am tempted to dismiss it as a commonplace; but when I remember other applications of linguistics (or other disciplines, such as computing) to literature I am reminded that though it _should_ be a commonplace it decidedly is not. As I said at the beginning of this review, an intelligent, sensitive reader of poetry is sharing insights with us. And there is no need to defend every line to yet think it excellent.

CURRENT BIBLIOGRAPHY

AJCL thanks Martin and Iris Kay and Xerox Palo Alto Research Center for their help in preparing this bibliography.

Controversy over *Computer Power and Human Reason* - I

Benjamin Kuipers
MIT AI Laboratory, Cambridge, Massachusetts 02139

John McCarthy
Stanford University AI Laboratory, California 94305

Joseph Weizenbaum
Laboratory for Computer Science, MIT, Cambridge, Massachusetts 02139

SIGART Newsletter 58: 4-12, June 1976

Kuipers: *Reactions to Weizenbaum's book*, 4-5. McCarthy: *An Unreasonable Book*, 5-10. Weizenbaum: *A Response to John McCarthy*, 10-12. Kuiper's reactions are fairly moderate in criticism of Weizenbaum while McCarthy's objections, given in some detail, are aimed deeper. Weizenbaum says: "The distance that separates John McCarthy from Joseph Weizenbaum is truly measured by the challenges these two hurl at one another: McCarthy defies Weizenbaum to 'Show me a way to knowledge besides science!' And Weizenbaum responds: 'Can there be a way toward an authentic model of man that does not include and ultimately rest on philosophical and moral thinking?'"

Controversy over *Computer Power and Human Reason* - II

Roger C. Schank
Computer Science Department, Yale University, New Haven, Connecticut

Kenneth Mark Colby
Algorithmic Laboratory of Higher Mental Functions, UCLA Department of Psychiatry, Los Angeles, California 90024

Josephy Weizenbaum
Laboratory of Computer Science, MIT, Cambridge, Massachusetts 02139

SIGART Newsletter 59: 7-11, August 1976

Schank: *Response to 'Weizenbaum's Response to McCarthy*, 7-8. Weizenbaum: *Reply to Roger Schank*, 8-9. Colby: *On the Morality of Computers Providing Psychotherapy*, 9-10. Weizenbaum: *Response to Colby*, 10-11. Schank and Weizenbaum are concerned with epistemological status of linguistic theories constructed by linguists (Chomsky in particular) and NL 'understanding' programs constructed by AI researchers, (e.g. Schank). Colby and Weizenbaum discuss the (im)morality of computer administered psychotherapy.

Computer Science as Empirical Inquiry

Allen Newell, and Herbert A. Simon
Carnegie-Mellon University, Pittsburgh, PA

Communications of the ACM 19: 113-126, March 1976

Two hypotheses concerning the nature of intelligence and AI are advanced and explicated: 1) *The Physical Symbol System Hypothesis.* A physical symbol system has the necessary and sufficient means for general intelligent action. 2) *Heuristic Search Hypothesis.* The solutions to problems are represented as symbol structures. A physical symbol system exercises its intelligence in problem solving by search - that is, by generating and progressively modifying symbol structures until it produces a solution structure.

GENERAL

Machine Understanding

Dennis de Champeaux
Instituut voor Bedrijfsinformatics, University of Amsterdam, Amsterdam-1001, Netherlands

April 1975

A brief overview of current NL programs, with most attention on the process of translating input sentences to a meaning representation. The impossiblity of delineating independent sub-problems in this process is discussed. Still, a rank ordering is given of *simple* representations of the following sub-problems: word recognition, syntax analysis, establishing existence of references, function of sentence parts, word ambiguity, case ambiguity, reference ambiguity, sentence type recognition, role-theme-scenario-etc.- recognition. Examples are given however of these sub-problems which can be solved only by first (partially) solving sub-problems which rank higher.

Dreyfus's Disproofs

Yorick Wilks
Department of Computation Logic, University of Edinburgh, EH8 9NW

British Journal for the Philosophy of Science 27: 177-185, June 1976

Dreyfus' 1972 book *What Computers Can't do* is open to 4 major lines of criticism: 1) Dreyfus doesn't describe the AI enterprise correctly, 2) AI work has changed considerably since his survey and is beginning to incorporate that which he claimed an adequate AI system must, 3) there is a lack of clarity in his claims, especially in the notion of "digital" (as opposed to analog) machines, 4) his phenomenological argument shows that machines cannot understand, but at the old philosophical price of showing that no one but I can, either (solipcism).

Uses of Higher Level Knowledge in a Speech Understanding System

William A. Woods, Madeleine Bates, Bertram Bruce, and Bonnie Nash-Webber
Bolt Beranek and Newman, Inc., Cambridge, Massachusetts 02138

Continuous speech understanding, for all but the most restricted languages, requires a characterization of the sentences of the language, the meanings of the words within the domain, and the use of the language within the task. Thus, syntax, semantics and pragmatics become essential components of a total speech understanding system. The control problem in such a system involves integrating the knowledge available from these components with the possible interpretations of an utterance suggested by its acoustics. This paper discusses various strategies for controlling these high-level sources of knowledge, reflecting their potential for coping with the inevitable acoustic ambiguity and error.

GENERAL

The Prague Bulletin of Mathematical Linguistics 25

Universita Karlova, Praha, 1976

CONTENTS

GENERAL

The Prague Bulletin of Mathematical Linguistics 26

Universita Karlova, Praha, 1976

CONTENTS

Annual Bulletin of the Research Institute of Logopedics and Phoniatrics, No. 10

Faculty of Medicine, University of Tokyo, March 1976

CONTENTS

*These articles have been abstracted elsewhere on this fiche.

Acoustic-Phonetic Experiment Facility for the Study of Continuous Speech

Richard M. Schwartz
Bolt Beranek and Newman, Inc., Cambridge, Massachusetts 02138

While gathering acoustic data for the acoustic-phonetic analysis of speech, it is necessary to consider many different sounds in varying phonetic environments to assure that the results are statistically significant. In order to reduce the amount of time required to test hypotheses, a facility has been developed which provides an interactive environment for performing a wide variety of acoustic-phonetic experiments on a large data base of continuous speech. Using this facility, one can formulate an experiment, run it on selected portions (or all) of the data base, and display or tabulate the results in a meaningful way. Another experiment may then be run based on the results. CPU time required to run an experiment on the entire data base is between 5 and 20 seconds, depending on the complexity of the experiment. Due to the ease of interactions, formulating or revising an experiment, running it, and displaying the results normally takes less than 5 minutes. This facility has been used in combination with a data base of 69 hand-labeled sentences to develop algorithms for acoustic-phonetic segmentation and labeling in a speech understanding system. Several examples of its use and the results obtained are presented.

PHONETICS-PHONOLOGY: PHONOLOGY

Analysis, Recognition, and Perception of Voiceless Fricative Consonants in Japanese

Hiroya Fujisaki
Research Institute of Logopedics and Phoniatrics, University of Tokyo

Osamu Kunisaki
Department of Electrical engineering, Faculty of Engineering, University of Tokyo

Annual Bulletin of the Research Institute of Logopedics and Phoniatrics 10: 145-156, 1976

The model for the spectral characteristics of voiceless fricative consonants in Japanese is based on an equivalent circuit representation of their generation mechanism. The model, together with its three simplified versions, is then evaluated from the point of view of automatic recognition as well as of synthesis of speech. For automatic recognition, spectral models that contain zeros are found to be particularly effective, and their parameters are shown to be sufficient for the complete separation of /s/ - and /sh/ - samples in CV and VCV utterances. On the other hand, perceptual experiments using synthetic stimuli reveal considerable smaller differences between models with spectral zeros and those without zeros.

Dictionary Expansion via Phonlogical Rules for a Speech Understanding System

William A. Woods, and Victor W. Zue
Bolt Beranek and Newman, Inc., Cambridge, Massachusetts 02138

This paper describes some modifications and extensions to the Bobrow-Fraser Phonological rule testing program (CACM 11, 1968, 766-772), which enable it to expand a dictionary of words read from a file and systematically try all possible combinations of optional rules, compute a running "probability" of a pronunciation as a function of probabilities of application and non-application of individual optional rules, associate arbitrary applicability tests with rules, and generate a summary table of which rules applied to each word and which words a given rule applied to. The system uses a notation similar to the classical Chomsky-Halle notation for phonological rules, and has been used to expand a 500-word dictionary for the BBN SPEECHLIS speech understanding system.

PHONETICS-PHONOLOGY: RECOGNITION

Methods for Nonlinear Spectral Distortion of Speech Signals

John Makhoul
Bolt Beranek and Newman, Inc., Cambridge, Massachusetts 02138

The spectral distortion of speech signals, without affecting the pitch or the speed of the signal, has met with some difficulty due to the need for pitch extraction. This paper presents a general analysis-synthesis scheme for the arbitrary spectral distortion of speech signals without the need for pitch extraction. Linear predictive warping, cepstral warping, and autocorrelation warping, are given as examples of the general scheme. Applications include the unscrambling of helium speech, spectral comprassion for the hard of hearing, bit rate reduction in speech compression systems, and efficiency of spectral representation for speech recognition systems.

Towards Perceptually Consistent Measures of Spectral Distance

Viswanathan, R. (Vishu), J. Makhoul, and W. Russell
Bolt Beranek and Newman, Inc., Cambridge, Massachusetts 02138

This paper considers distance measures for the determining the deviation between two smoothed short-time speech spectra. Since such distance measures are employed in speech processing applications that either involve or relate to human perceptual judgment, the effectiveness of these measures will be enhanced if they provide results consistent with human speech perception. As a first step, we suggest Flanagan's results on difference limens for formant frequencies as one basis for checking the perceptual consistency of a measure. A general necessary condition for perceptual consistency is derived for a class of spectral distance measures. A class of perceptually consistent measures obtained through experimental investigations is then described, and results obtained using one such measure under Flanagan's test conditions are presented.

PHONETICS-PHONOLOGY: RECOGNITION

Parametric Matching in a Word Verification Component

Craig C. Cook
Bolt Beranek and Newman, Inc., Cambridge, Massachusetts 02138

One of the components of the BBN Speech Understanding System is a Word Verification component. Composed of a synthesis-by-rule program and a parametric word verifier, this component useds the synthetic acoustic representation of a hypothesized word to check for the presence of that word at a particular location in an unknown utterance. We describe the strategies of the verifier which includes a time normalization algorithm and a spectral matching error metric. The application of these techniques to continuous speech understanding in a multi-speaker environment is also discussed.

Temporal Organization of Articulatory and Phonatory Controls in Realization of Japanese Word Accents

Hiroya Fujisaki
Research Institute of Logopedics and Phoniatrics, Faculty of Medicine, University of Tokyo

Hiroyoshi Morikawa
Department of Electrical Engineering, Faculty of Engineering, University of Tokyo

Miyoko Sugito
Department of Japanese Literature, Faculty of Education and Liberal Arts, Osaka Shoin Women's College

Annual Bulletin of the Research Institute of Logopedics and Phoniatrics 10:177-190, 1976

In the analysis of word accent in Japanese segmental feactures were extracted as formant frequency trajectories from short-time frequency spectra of speech, while suprasegmental features were extracted as the contour of the fundamental frequency of the glottal source. The temporal relationships between these features were quantified by analyzing the time-varying patterns of these acoustic parameters on the basis of functional models for their control processes, and the instants of initiation of articulatory and phonatory controls were estimated, the results indicating that they were not exactly simultaneous but were possibly brought into an approximate synchronism by a mediating process. The possible role of perception as the mediating process was then examined both by identification of truncated utterances and by click location to determine perceptual segment boundaries. It was found that the perceptual segment boundary roughly concided with the instant of initiation of the phonatory control characteristic of each word accent type, thereby suggesting the role of perception in the coordination of articulatory and phonatory controls.

New Lattice Methods for Linear Prediction

John Makhoul
Bolt Beranek and Newman, Inc., Cambridge, Massachusetts 02138

This paper presents a new formulation for linear prediction, which we call the covariance lattice method. The method is viewed as one of a class of lattice methods which guarantee the stability of the all-pole filter, with or without windowing of the signal, with finite wordlength computations, and with the number of computations being comparable to the traditional autocorrelation and covariance methods. In addition, quantization of the reflection coefficients can be accomplished within the recursion for retention of accuracy in representation.

PHONETICS-PHONOLOGY: RECOGNITION: SYSTEM

Acoustic-Phonetic Recognition in BBN SPEECHLIS

Richard M. Schwartz, and Victor W. Zue
Bolt Beranek and Newman, Inc., Cambridge, Massachusetts 02138

The system accepts various parameters derived from the digital waveform and short-time spectra, and produces a segment lattice where segments can have overlapping boundaries and the description of segments is a list of labels. Acoustic-Phonetic as well as phonological knowledge of English is employed extensively in labeling the segments. Each label also has associated with it a score, reflecting the confidence in its identity. A description of the acoustic-phonetic analyzer, as well as statistics related to its performance are presented in detail.

Analysis, Synthesis, and Perception of Word Accent Types in Japanese

Hiroya Fujisaki, and Hajime Hirose
Research Institute of Logopedics and Phoniatrics, Faculty of Medicine, University of Tokyo

Miyoko Sugito
Department of Japanese Literature, Faculty of Education and Liberal Arts, Osaka Shoin Women's College

Annual Bulletin of the Research Institute of Logopedics and Phoniatrics 10:173-176, 1976

Fundamental frequency of the glottal source (F_0) is the primary acoustic correlate of prosodic features such as word accent and intonation in spoken Japanese. In order to separate linguistically relevant information from characteristics of the glottal mechanisms of glottal oscillations that generates the F_0-contour from binary commands of voicing and accent. This paper concerns the extension of the model from the *Tokyo* dialect to *Kinki* dialects, such as *Osaka* and *Kyoto*, which possess accent types which are marked with pitch transitions within a mora. Experiments indicate that perception of a particular accent type is closely tied to a specific temporal relationship between a segment boundary and the onset or the offset of the accent command.

Acoustic and Perceptual Analysis of Two-Mora Word Accent Types in the *Osaka* Dialect

Hiroya Fujisaki
Research Institute of Logopedics and Phoniatrics, Faculty of Medicine, University of Tokyo

Miyoko Sugito
Department of Japanese Literature, Faculty of Education and Liberal Arts, Osaka Shoin Women's College

Annual Bulletin of the Research Institute of Logopedics and Phoniatrics 10: 157-171, 1976

In order to investigate the acoustic characteristics which play a major role in the generation as well as in the perception of word accent types in Japanese, pitch contours of all accent types found in two-mora words of the *Osaka* dialect have been analyzed. Distinctions between accent types have been found to be most clearly represented by the timing of the onset and offset of the underlying accent command extracted from a measured pitch contour. Perceptual importance of these parameters has been confirmed by identification tests using synthetic speech stimuli with a variety of pitch contours. Results of further perceptual experiments have revealed that the identification of an accent type is based predominantly on the relative timing of the accent command and specific segment boundary within a word.

The Use of Context in Character Recognition

Edward G. Fisher
Department of Computer and Information Science, University of Massachusetts, Amherst 01002

COINS Technical Report 76-12

Improvements to binary n-gram algorithms are proposed which facilitate the use of different types of n-grams for a collection of words of varying lengths by using an extended data base for the detection, location, and correction of insertion, deletion, split, and merger errors. A primary feature of the new algorithms is that the n-grams are anchored to one or both ends of the word. Experiments show that the algorithms are effective for all of the error types. For example, in a dictionary of 10,000 words the contextual post-processor was able, without *a priori* knowledge of the types of errors being processed, to correct 81% of single substitution errors, 57.7% of deletions, and 63.2% of double substitution errors. Application of these techniques to postal reading machines is explored.

WRITING: RECOGNITION: CHINESE.

Machine Processing of Chinese Characters

William Stallings
Center for Naval Analysis, Arlington, VA

SIGLASH Newsletter 9, No. 3: 9-11, June 1976

The large number of characters, and their complexity, makes processing of Chinese extremely difficult. Keyboard entry is slow, with high error rates, and a great deal of training is required to operate any of the various keyboards. The alternative input device - OCR - is not very promising as the large number of complex characters makes discrimination difficult.

What is the Proper Characterization of the Alphabet, I. Desiderata

W. C. Watt
Department of Cognitive Sciences, University of California, Irvine 92664

Visible Language 9: 293-327, Autumn 1975

In their characterization of the alphabet, Eden-Halle (1961) attacked the right domain (their phoneme-analogues were strokes, not letters), while Gibson et al. (1965, 1969, 1963) came closer to employing the right kind of distinctive features (theirs were much more freely combinable and had some psychlogical warrant). The obvious next move should be something like: the printed majuscules should be dissected into their proper strokes - as Eden and Halle did for the cursives - and those strokes should then be analyzed in terms of Gibson-like distinctive features ("verticle," "curved," etc.). Global features, such as "symmetric," can be derived from sequences of line-segments expressed without them.

LEXICOGRAPHY-LEXICOLOGY: THESAURI

A Classification of the "matter concepts" for Linguistic Recognition Picture Patterns (In Japanese)

Kyushu daigaku kogaku syuho., Technological Reports of the Kyushu University 46: 560-569, 1973

A file of 5400 Japanese verbs has been used for one-dimensional classification by structural and semantic features. The classification was based on a consecutive dichotomic division of verbs into groups.(verbs of "action" -verbs of "non-action", "simple verbs"- compound verbs, etc.); the absolute synonyms were then reduced to equivalence classes (for instance, *korogara=korogera=korobu=marobu=*to fall, to roll); the final stage of the classification brought the verbs together into semantic classes and isolated "the basic concepts". Altogether 1200 "basic concepts" have been singled out. The compound verbs are analyzed in terms of their morphological structures, taking account of the potential verbal government. The semantic features of verbs as written in the machine dictionary are intended for recognition of various linguistic objects, in particular for identifications are to be used in computer experiments. 5 refs.

A Bi-Directional Speech Parsing Technique

Carol C. Reason
Department of Computer Science, University of Toronto, M5S 1A7, Canada

Technical Report 90

This method of bi-directional parsing of continuous speech involves data objects (Assemblies) which are similar to single paths in a ATN. The objects specify the language grammar and at the same time are executed as the parsing mechanism and can make use of non-phonetic data in the speech signal, which may help disambiguate the parse of various language constructs. Assemblies that are partially or completely realized in the signal become theories about the utterance context, which in turn may be used by other Assemblies. The control structure is capable of producing more than one parse for a particular utterance and so leaves the decision as to which is 'correct' to some higher level process.

GRAMMAR: PARSER

Comprehension by Computer: Expectation-Based Analysis of Sentences in Context

Christopher K. Riesbeck, and Roger C. Schank
Department of Computer Science, Yale University, New Haven, Connecticut

Research Report 78, 76 pp, October 1976

ELI (English Language Interpreter) maps English sentences into their conceptual representation (in Conceptual Dependency form) and is not a parser in the ordinary sense of the word; for it lacks an intermediate syntactic analysis stage. ELI is implemented as part of SAM (Script Applier Mechanism) and interacts with other parts of SAM a great deal. The following principles have been followed in designing ELI: 1) A parse must be done in (semantic) context only. Context eliminates ambiguity, which leads to the second principle. 2) A parser should never notice ambiguity. 3) A parser must take care of syntactic consiserations only when required to do so by semantic considerations. 4) Parsing is expectation-based. 5) Words with multiple senses are ordered with respect to context. 6) Parsing is really a memory process. The basic memory process of testing for conditions and making predictions based on the results of those tests is fundamental to both parsing and higher level processes.

Augmented Phrase Structure Grammars

George E. Heidorn
IBM Thomas J. Watson Research Center, Yorktown Heights, New York 10598

Research Report RC 5787, 15pp, December 31, 1975

An augmented structure grammar (APSG) consists of a collection of phrase structure rules which are augmented by arbitrary conditions and structure building actions. APSG uses a data structure in the form of a semantic network consisting of 'records' which are collections of attribute-value pairs. The construction of such records from strings of text proceeds by decoding rules each of which has a list of one or more 'segment types' (meta-symbols) on the left of the arrow to indicate which types of contiguous segments must be present in order for a segment of the type on the right of the arrow to be formed. Applicability condiditons may be stated in parentheses on the left side or the rule and structure building operations in the creation of a new segment are stated in parentheses on the right side. For encoding, the right side of the rule specifies what segments a segment of the type on the left side is to be expanded into. Conditions and structure-building actions are included in exactly the same manner as in decoding rules.

GRAMMAR: PARSER

An Algebraic Technique for Context-Sensitive Parsing

Robert L. Cannon, Jr.
Department of Mathematics and Computer Science, University of South Carolina, Columbia

International Journal of Computer and Information Sciences 5: 257-276, September 1976

A technique that represents derivations of a context-free grammar *C* over a semiring and that obtains for a word *w* in *L(G)* the set of all canonical parses for *w* has previously been described. A state grammar is one of a collection of grammars that place restrictions on the manner of application of context-free-like productions and that geneate non-context-free language. Weiss, Nago, and Stanat (1973) have developed a technique for state grammar parsing where an algebraic expression is derived from the context-free grammar and the input string. For each possible sequence of derivations from the input string, a term yielding that derivation sequence appears in the expression. This powerful formal tool has been extended for state grammars by utilizing as much as possible the properties of a state grammar that are derived from its context-free origins.

Parsing of General Context-Free Languages

Susan L. Graham, and Michael A. Harrison
Computer Science Division, University of California at Berkeley

Morris Rubinoff and Marshall Yovits, eds., Advances in Computers, 14 (Academic Press ISBN 0-12-012114-X): 77-185, 1976

The Cocke-Kasami-Younger Algorithm: the recognition algorithm, the parsing algorithm, a Turing machine implementation, linear grammars - a special case. *Earley's Algorithm:* the recognition algorithm, correctness of the algorithm, the time and space bounds, the parsing algorithm. *Valiant's Algorithm:* recognition as a transitive closure problem, Strassen's algorithm and Boolean matrix multiplication, Valiant's lemma, computing D+ in less than $O(n^3)$ time, an upper bound for context-free parsing. *The Hardest Context-Free Language. Bounds on Time and Space.*

A Process to Implement Some Word Sense Disambiguations

Phillip Hayes
Computer Science Department, University of rochester, New York 14627

Technical Report 6

Word-sense ambiguity is one of the major problems to be faced by any natural language understanding system. A process is proposed to implement disambiguations, explainable by four classes of influence -- selectional restrictions, preference restrictions, logical restrictions, and association. The process operates as an interface between parsing and deeper processing. The treatment of association is based on a frame-like organizational system for semantic nets. This system provides a means of representing context and facilitates the search for an associative connection between a word-sense and an already existing context. In addition, it is argued that selectional restrictions really are absolute, even though sentences exist in which they are apparently violated.

Knowledge in Automatic Planning Systems

Richard E. Fikes
Stanford Research Institute, Menlo Park, CA 94025

Technical Note 119 (Pub. Z155, $2.00)

A tutorial on automatic planning systems with particular emphasis on knowledge representation issues. The effectiveness of a planner depends on its ability to make use of descriptions and expertise associated with particular task domains (such as action models, state description models, scenarios, and special purpose plan composition methods). The discussion focuses on how such knowledge is represented in planning systems and how various planning strategies.

Computational Semantics: An Introduction to Artificial Intelligence and Natural Language Comprehension

Charniak Eugene, Ed.
Institute for Semantic and Cognitive Studies, University of Edinburgh

Fundamental Studies in Computer Science, Volume 4, Amsterdam, NewYork, Oxford: North-Holland Publishing Company, 1976
xiii + 294 pp, Dfl. 55.00/$19.00, ISBN 0-7204-0469-X

CONTENTS

SEMANTICS-DISCOURSE: THEORY

Extending the Expressive Power of Semantic Networks

L. K. Schubert
University of Alberta, Edmonton, Canada

Artificial Intelligence 7: 163-198, Summer 1976

A network representation is developed which permits the use of *n*-ary predicates, logical, connectives, unrestricted quantification (including quantification over predicates), lambda abstraction, and modal operators such as belief and counterfactual implication. The representation easily accommodates propositions of the type encoded by Quillian, Winston, Schank, and Rumelhart *et al.* in their networks. All extensions of the basic propositional notation are analogues of standard notational devices employed in various first-order or higher-order Predicate Calculi.

SEMANTICS-DISCOURSE: THEORY

Semantic Interpretation Revisited

Bonnie Nash-Webber
Bolt Beranek and Newman, Inc., Cambridge, Massachusetts 02138,

A brief overview is given of the BBN LUNAR system. This is followed by a discussion of two of its deficiencies: its purely extensional processing of its meaning representation languate and its inadequate treatment of anaphora. We then present a rough classification of anaphoric expressions as groundwork towards formulating a general computational treatment of the phenomenon. In this classification, we establish a distinction between *denotational anaphora* -- references to previously mentioned objects, sets, events, states, etc. -- and *descriptional anaphora* -- references to previous descriptions. Finally, we present an initial sketch of both a formal meaning representation language and some procedures seen needed for manipulating sentences of that language, which may provide a handle on some aspects of anaphor resolution.

SEMANTICS-DISCOURSE: THEORY

How Near is Near?

Murray Elias Denofsky
Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge 02139

Memo 344, 75 pp., $2.10, February 1976

This paper presents a system for understanding the concept of 'near' and 'far', weighing such factors as purpose of judgment, dimensions of the objects, absolute size of the distance, and size of distance relative ot other objects, ranges, and standards. A further section discusses the meaning of phrases such as 'very near', 'much nearer than', and 'as near as'. Although we will speak of 'near' as a judgment about physical distance, most of the ideas developed will be applicable to any continuous, measurable parameter, such as size or time. An adaptation for rows (discrete spaces) is made as well.

A Model of Concept Formation

D. J. H. Brown
Applied Mathematics Department, University of the Witwatersrand, Johannesburg

Preprint No. A3 - 1974

A program capable of forming concepts and generating English descriptions of them is described. Concepts are formed from advice and/or positive and negative exemplars. The memory takes the form of a context-dependent net of productions, upon which introspection can be performed. The program has been applied to several domains and exhibits primitive learning capabilities. The program is written in Algol 68-R and LISP 1.5.

Concept Formation and Exposition

D. J. H. Brown
Applied Mathematics Department, University of the Witwatersrand, Johannesburg

Ph.D. Dissertation, July 1975

The model of concept formation utilizes a memory structure in the form of a context-dependent network of production rules. Techniques for the induction of such rules are examined and discussed. Heuristic measures are introduced which permit the model to introspect upon its knowledge. Additionally, the concept of a Generalized Extended-Entry Decision Table (GEEDT) is presented, and a method of transforming a memory structure into a GEEDT is examined. Transformations from GEEDT's into decision trees and English Statements are Described. Finally the behavior of the model is examined by applying it to four different tasks.

Qualitative and Quantitative Knowledge in Classical Mechanics

Johan de Kleer
Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge 02139

AI Technical Report 352, 113 pp., $2.00, December 1975

A program called NEWTON understands and solves problems in a mechanics mini world of objects moving on surfaces. Facts and equations such as those given in a mechanics text need to be represented, but they are not sufficient to solve problems. Human problem solvers also rely on qualitative knowledge which the physics text tacitly assumes. The representation of both these kinds of knowledge is the major preoccupation of this work. The knowledge is used by a planning problem solver. Planning involves tentatively outlining a possible path to the solution without actually solving the problem. "Envisioning," or qualitative simulation of events, plays a central role in this planning process.

SEMANTICS-DISCOURSE: COMPREHENSION

Computer Understanding of Physics Problems Stated in Natural Language

Gorden S. Novak, Jr.
Computer Science Department University of Texas, Austin 78712

Ph. D. Thesis, 1976

ISAAC is a program which can read, understand, solve, and draw pictures of physics problems stated in English. The problems involve rigid bodies in static equilibrium, and include such objects as levers, pivots, weights, ropes, and springs in various configurations. The understanding of a physics problem is an active process in which the sentences of the problem statement are used to guide the construction of a model which represents the relationships and features of objects with much greater detail and specificity than is present in the original problem statement. These representations are further elaborated by processes which construct a geometric model of the problem, associate canonical objects (such as a point mass) with physical objects (such as a person), write and solve equations which describe the interactions of the objects, and construct a diagram of the problem.

An Application-Independent Subsystem for Free-Text Analysis

Robert R. Fenichel
Department of Internal Medicine, Center for the Health Sciences, University of California at Los Angeles, 90024

G. Octo Barnett
Laboratory of Computer Science, Massachusetts General Hospital, Boston, 02114

Computers and Biomedical Research 9: 159-167, April 1976

The analysis of free English text is a major subproblem in medical computing. A free-text-analysis subsystem, written entirely in MUMPS language and built around a numerical taxonomy of anticipated input services multiple application programs at Massachusetts General. Facilities are provided for matching, spelling correction, disinflection, and phrase analysis. A variety of administrative programs are provided for logging uninterpreted input and for dynamically incrementing the data-base so as to expand the range of correctly interpreted input.

SEMANTICS-DISCOURSE: MEMORY: QUESTION ANSWERING

A Question-Answering System (QUANSY) for Information Retrieval

Abraham S. Ben David
Department of Information Science, Lehigh University, Bethlehem, Pensylvania 18015

Ph.D. dissertation, 161 pp., 1976

QUANSY 3.0 has the processing power of other current 'state of the art' question answering systems and is faster and more flexible than most. As there seems to be no satisfactory overall theory for NL question answering systems QUANSY has been developed without the use of an overall unifying theory. The development of QUANSY proceeded by testing every theory and idea deemed relevant, incorporating those that work, discarding those that don't; more work of this sort is necessary before development of unified theories becomes feasible. QUANSY has three main components: the *memory* component receives information from the *parser-analyzer* and the *question-answering routines* and disseminates information to them (in effect, they take what they need when they need it). The system uses various types of inference in answering questions posed in NL. Specifics of the memory structure, the grammatical analysis, and the question answering routines are given. The appendices include a sample dialog and a discussion of the dialog.

SEMANTICS-DISCOURSE: TEXT GRAMMAR

Experiments in Conceptual Analysis of Theoretical Discourses

Jean-Guy Meunier
Universite du Quebec a Montreal, Case postale 8888, H3C 3P8, Canada

Paper presented and the International Conference on Computational Linguistics, Ottawa, 1976

The process to be modeled in text analysis is the extraction of information and meaning which is new to the system from an utterance it has never before encountered. One of the primary problems in investigating literary or philosophical texts is their use of terms, such as life, soul, freedom, etc., which do not permit definition in terms of physical entities or even simple mental events. Nor does the meaning of these concepts necessarily remain stable throughout the course of a given text; often the word is being defined in the unfolding of the discourse. The study of the lexicon is seen as an entrance door to the study of the text. Each word receives its final meaning when it actually occurs in the text as a whole, and the text is slowly built by the semantic transformation of each of its constituents. Currently three methods are being used in the investigation of a text from Descartes *Discours de la method.* The first method is based in information retrieval methods and various statistical techniques (stars, clic, chain, clumps, clumps representation, cluster and factor analysis). Some results of an analysis of *connaitre, penser,* and *savoir* are given. The second approach is syntactic. The third approach is semantic and has not bet been implemented. The approach is modeled on Wilks semantic preference and uses 85 primitives combined according to a set of 25 rules.

SEMANTICS-DISCOURSE: CONCEPTUAL DEPENDENCY

An Organization of Knowledge for Problem Solving and Language Comprehension

Chuck Rieger
Department of Computer Science, University of Maryland, College Park 20742

Artificial Intelligence 7: 89-127, Summer 1976

Commonsense algorithm representations (CARs) can be built from 26 connective links between 5 types of actions, two of which are states and statechanges: actorless conditions in the world which can be caused by actions. A bypassable causal selection network is a tree where each node has a test and a set of one or more alternative branches associated with it. One such network exists for each state or statechange concept. The terminal nodes of these networks contain *approaches,* CARs which map out plans for solving problems. All networks tests are to be shared; there is a central copy of the test and references to it from each point in the net where the test is to be applied. By making appropriate use of this feature it is possible to insert conditional bypasses into the network such that, as more and more of the environment becomes known the solutions to problems tend to grow increasingly stereotyped. Thus bypassable networks can adapt to context. The bypass arrangement is frame-like, but the frames are distributed bundles of bypasses which can blend in essentially infinite variety. Such a network can be used by a plan synthesizer and a language comprehender.

A Conceptual Approach to Automated Language Understanding and Belief Structures: With Disambiguation of the Word 'For'

Linda Gail Hemphill
Stanford University AI Laboratory, CA 94305

Ph.D. Thesis in Linguistics; AIM-273, STAN-CS-75-534, 254pp, May 1975

The word 'for' is given special attention in this thesis about language understanding because it can have more than twenty different meanings, and yet, a person rarely misinterprets an instance of it or finds it ambiguous. The model must make inferences from the sentences under analysis; it must analyze two syntactically different sentences which are paraphrases of each other into the same semantic representation; and it must interact with a memory structure. The types of information needed for an understanding system to correctly interpret 'for' are shown to interact in other instances of language understanding and generation as well.

SEMANTICS-DISCOURSE: EXPRESSION

The Metanovel: Writing Stories by Computer

James Richard Meehan
Department of Computer Science, University of California, Irvine

Department of Computer Science, Yale University, Research Report 74, 237 pp, September 1976

TALE-SPIN produces stories in English in one of two general modes. In the bottom-up mode the program interacts with the user, who specifies characters, personality characteristics, and relationships between characters. In the top-down mode it produces Aesop-like fables. To do this it requires that the program user specify some moral (such as 'never trust flatterers,' 'don't take a dare,' etc.) and it then produces an appropriate story with the system itself making decisions about the story which had been up to the user in bottom-up mode. The system uses Conceptual Dependency (Schank) structures to model world knowledge required in story builging: knowledge of the physical world; rules of social behavior and relationships; techniques for solving everyday problems such as transportation, acquisition of objects, and acquisition of information; knowledge about physical needs such as hunger and thirst; knowledge about planning behavior and the relationships between kinds of goals; and knowledge about stories, their organization and contents and how to express them in English.

The Translation of Formal Proofs into English

Dahiel Chester
Department of Computer Sciences, University of Texas, Austin 78712

Artificial Intelligence 7: 261-278, Fall 1976

EXPOUND translates a formal proof into an English statement of a theorem and its proof. In the first stage it makes a graph representation of the inferential relationships between the lines of the proof. In the next two stages it uses this graph to make an outline of the text which it will generate. The program makes this outline by first grouping lines together into paragraphs, then putting these paragraphs in linear order and inserting introductory paragraphs where appropriate to explain how the other paragraphs are related. Finally, it generates an English text by explaining how each line is obtained from the preceding lines in the outline. Some lines in the formal proof are not translated since the inferences deriving them are trivial and easily reconstructed by the reader. The grammar used is a simple case grammar; each predicate has assoicated with it a verb string, a syntax type, and the preposition, if any, associated with each argument.

LINGUISTICS: METHODS: MATHEMATICAL

Grammars of Partial Graphs

H. J. Schneider
Institut fuer Math. Maschinen und Datenverarbeitung (II) der Univ. Erlangen-Nuernberg, D-8520 Erlangen Federal Republic of Germany

H. Ehrig
Forschungagruppe Automatentheorie und Formale Sprachen der TU Berlin, Otto-Suhr-Allee 18-20, D-1000 Berlin 10

Acta Infromatica 6: 297-316, 1976

The concept of Chomsky-grammars is generalized to graph-grammers; the "gluing" of graphs is defined by a pushout-construction. We allow the left-hand and right-hand side of a production to be partial graphs, i.e. graphs in which there may be edges without a source or target node. A necessary and sufficient condition for applicability of productions is given and convex graph-grammars are studied.

Prospects for Parallelism and the Computer Crunch

John M. Carroll
Computer Science Department, University of Western Ontario, London N6A 3K7, Canada

Journal of the American Society for Informantion Science 27(1): 63-69, January-February 1976

The possibilities of parallel processing are explored with respect to three functions fundamental to information science: sorting, searching, and evaluation of a search prescription In a simulation experiment involving sorting and exchange sort in 10 simulated parallel registers combined by a decommutator model wasn't quicker than a sort using the Shell or Quicksort algorithms. This result was predicted and more' efficient parallel sort strategies are suggested (but haven't yet been subjected to test by simulation). Four string search strategies were tested, the simulation of parallel processing being the fastest of the four.

COMPUTATION: INFERENCE

Generalized AND/OR Graphs

Giorgio Levi, and Franco Sirovich
Instituto di Elaborazione della Informazione, Pisa, Italy

Artificial Intelligence 7: 243-259, Fall 1976

A Generalized AND/OR Graph (GAG) is defined as $G = (O, T, A, E, S)$, where O is a set of *nonterminal OR nodes* (problems), X T is a set of *terminal OR nodes* (solved problems), A is a set of *AND nodes* (problem reduction operators), E is a subset of (O U T) X A U A X (O U T) whose elements are directed arcs, S belongs to O and is called the *start node*. In a GAG reduction operators can apply to more than a single input problem, which allows a GAG to deal with subproblem interdependence in problem reduction. An ordered-search algorithm, proven to be admissible and optimal, if given to find a solution. Examples are given which show the application of the formalism to problems which cannot be modelled by AND/OR graphs. GAGs are shown to be equivalent to type O grammars so that finding a solution of a GAG is equivalent to deriving a sentence in the corresponding type O grammar.

Natural Resolution: A Human-Oriented Logic Based on the Resolution Principle

Daniel J. Buehrer
Computer Science Department, University of Iowa, Iowa City 52242

Unlike most other refinements of resolution, natural resolution corresponds closely to other human-oriented problem-solving methodologies; problem-reduction searches, state-space searches, backtracking heuristics, goal statement tree searches, evaluation functions, and pattern-directed nonprocedural programs all have straightforward interpretations in terms of natural resolution (NR). NR makes use of a partially defined 'real-world' model which is consistent with the axioms. The model is used to classify phrases as either operands or operators. Operands are set-of-support phrases which have a true interpretation in the model. Operators are allowed to resolve on their false literals with the operands. NR uses a simple version of inductive logic to assign default likelihoods which are used to locate 'relevant' operators. The user may define his own likelihood functions to provide the program special hints for choosing operators. The concept of affected literals' is introduced into independent subproblems while dummy literals' are used to identify dependent subproblems. Othel human-oriented features of NR include data structed constants, variables with data types, and functions and literals which are automatically evaluated when they are defined as LISP functions. NR has been implemented in a program which has been tested on the same set of problems as QS3.6; results indicate that NR's use of a 'real-world' model does not significantly increase the search space.

COMPUTATION: PROGRAMMING

Four PL/1 Subroutines for Natural Language Processing

John Fink
Department of English, Rhode Island College, Providence

SIGLASH Newsletter 9, No. 3: 12-19, June 1976

1) FINDONE isolates words and punctuation marks in the input string of characters. 2) LAGADO1 is a word look-up function to be used with alphabetized type lists and uses the binary search technique. 3) LAGADO2 words much as LAGADO1 does except it works from a random order type list with has associated with it an index, array that alphabetizes the file. 4) LAGADO3 is a linear search look-up function. PL/1 code for the subroutines is given.

LAMBDA: The Ultimate Imperative

Guy Lewis Steele, Jr., and Gerald Jay Sussman
Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge 02139

AI Memo No. 353, 39 pp., $1.70, March 1976

We demonstrate how to model the following common programming constructs in terms of an applicative-order language similar to LISP: simple recursion, iterations, compound statements and expressions, GOTO and assignment, continuation-passing, escape expressions, fluid variables, call by name, call by need, and call by reference. The models require only lambda application, conditionals, and an occasional assignment. No complex data structures such as stacks are used. The models are transparent, involving only local syntactic transformations. Some of these models, such as those for GOTO and assignment, are already well known, and appear in the work of Landin, Reynolds, and others. The models for escape expressions, fluid variables, and call by need with side effects are new. The paper is partly tutorial in intent, gathering all the models together for purposes of context.

COMPUTATION: PROGRAMMING: LANGUAGES

QLISP; A Language for the Interactive Development of Complex Systems

E. Sacerdoti, R. Fikes, R. Reboh, D. Sagalowicz, R. Waldinger, and M. Wilber
Stanford Research Institute, Menlo Park, CA 94025

Technical Note 120 (Pub. Z156, $2.00), March 1976

QLISP is both a programming language and an interactive programming environment. It embeds an extended version of QA4, an earlier AI language, in INTERLISP The language features provided by QLISP include a variety of useful data types, an associative data base for the storage and retrieval of expressions, a powerful pattern matcher based on a unification algorithm, pattern-directed function invocation, "teams" of pattern involved functions, a sophisticated mechanism for breaking a data base into contexts, generators for associative data retrieval, and easy extensibility. See also *QLISP Reference Manual,* B. Michael Wilber, Technical Note 118 (Pub. Z154, $3.00), March 1976.

An INTERLISP Relational Data Base System

Stephen Weyl
Stanford Research Institute, Menlo Park, CA 94025

Technical Note 110 (Pub. Z152, $3.00), November 1975

INTERLISP has been augmented to support applications requiring large data bases maintained on secondary store. The data base support programs are separated into two levels: an advanced file system and relational data base management procedures which provide a high-level relational data base facility. Programmers define the logical contents of their data base as a set of relations declared as a formal data base definition, or "schema". This schema is then interpreted by general purpose functions that provide for creation, updating, and querying of relations.

DILOG - Digitalized Intelligence Logic Machine

H. K. Huan, and R. S. Ledley
National Biomedical Research Foundation, Georgetown University Medical Center, 3900 Reservior Rd., N.W., Washington, D.C. 20007

P. Johnson
Central Intelligence Agency

Computer Languages 2: 27-43, 1976

DILOG (Digitalized Intelligence Logic) is an interactive time-sharing computer language which is based on the principle of digital logical computational methods and built by using the APL (A Programming Language). This language is neither a general problem solver nor a model of reasoning; it is a man-machine interactive system for classical evaluation of statements by preprogrammed logical opertaions and algorithms. Four examples considered: 1) Mutually exclusive and exhaustive statements, 2) Medical diagnosis, 3) Analyzing experimental results (in biochemistry), 4) Analyzing intelligence reports.

COMPUTATION: PROGRAMMING: LANGUAGES

SCHEME: An Interpreter for Extended Lambda Calculus

Gerald Jay Sussman, and Guy Lewis Steele, Jr.
Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge 02139

AI Memo 349, 42 pp., $1.70, December 1975

Inspired by ACTORS, we have implemented an interpreter for a LISP-like language, SCHEME based on the lambda calculus, but extended for side effects, multiprocessing, and process synchronization. The purpose of this implementation is tutorial. We wish to: (1) alleviate the confusion caused by Micro-PLANNER, CONNIVER, etc. by clarifying the embedding of non-recursive structures in a recursive host language like LISP. (2) explain how to use these control structures, independent of such issues as pattern matching and data-base manipulations. (3) have a simple concrete experimental domain for certain issues of programming semantics and style. The paper includes a discussion of the issues facing an implementer of a language based on lambda calculus; and presents a completely annotated interpreter for SCHEME.

New Journal: *ACM Transactions on Database Systems*

David K. Hsiao, Editor-in-Chief
Department of Computer and Information Science, Ohio State University, Columbus 43210

Vol. 1, No. 1, March 1976

CONTENTS

Special Issue on Data-Base Management Systems

E. H. Sibley, Guest Editor

Computing Surveys 8, No. 1, March 1976

CONTENTS

COMPUTATION: INFORMATION STRUCTURES

Knowledge Space: A Conceptual Basis for the Organization of Knowledge

Peter P. M. Meinke
University of Toronto, Ontario M 5S 1A1, Canada

Pauline Atherton
School of Information Studies, Syracuse University, NY 13210

Journal of the American Society for Information Science 27(1): 18-24
January-February 1976

The organization of information can be conceptualized in terms of *concept 'vectors'* for a field of knowledge represented in a *multidimensional space*, and the *state 'vector'* for a person based on his understanding of these concepts, and the *representation 'vectors'* for information items which might be in a retrieval system which covers a subspace of knowledge. This accommodates the notion of *search volume* in which the user of a retrieval system can expand or reduce the subspace he searches for relevant information items which have representational vectors with components on basic concept vectors similar to his state vector. A system organized on this basis would require less information (on the part of the user) to describe and structure his search volume, and the state knowledge of the user could be monitered and recorded by the system with each transaction.

DATAPLAN: An Interface Generator for Database Semantics

Tosiyasn L. Kunii
IBM Research Laboratory, Computer Science Department, Monterey and Cottle Roads, San José, California 95193

Information Sciences 10: 279-298, 1976

A model system (DATAPLAN) is proposed which is capable of planning man-machine hierarchical dialogs to provide casual users with an interface to a shared database in their own terminology. The plan is executed on the domain of hierarchical semantics and controlled by pattern matching with the goal and uses a generative back-tracking mechanism. Adoption of the same representation for a database and interfacing data enables the causal user to control database contents by periodically merging both. The interface is individualized by division into user class files which, with multiple associations of assorted degree by "fuzzy" membership, ensures smooth communication of casual users with database. The values of the "fuzzy" membership functions are also periodically updated by the system based on the current database values.

Toward a Theory of Encoded Data Structures and Data Transitions

Ben Shneiderman, and Stuart C. Shapiro
Department of Computer Science, Indiana University, Bloomington, Indiana

Internation Journal of Computer and Information Science 5: 33-43, March 1976

Several models of data base systems have distinguished levels of abstraction ranging from the high-level entity set model down to the low-level physical device level. We present a model for describing data encodings, an intermediate level which focuses on the static relationships among data items as demonstrated by contiguity or by pointer connections. The model avoids issues related to physical devices and the details of pointer implementation and pursues a constructive approach which seeks to model the data fields and their contents, the data items. Multiple data encodings for a file are shown and transformation functions that describe the translation between data encodings are discussed.

A Selected Bibliography on Computer Vision

Martin D. Levine
Department of Electrical Engineering, McGill University, Montreal, Quebec, Canada

SIGART Newsletter 58: 17-19, June 1976

The bibliography contains 68 items arranged in the following categories: *Introduction*: 1.1 Introduction to AI, 1.2 Psychology and Vision, 1.3 Early Work. *Perception of Polyhedra*: 2.1 Polyhedra Recognition - Abstract, 2.2 Polyhedra Recognition - Real Data. *Natural Scene Analysis*: 3.1 Region Analysis, 3.2 Organizing Local Features; 3.3 Image Understanding. *Special Problems*: 4.1 Object Recognition, 4.2 Face Recognition, 4.3 Waveform Analysis. *Knowledge Representation*: 5.1 Image Descriptions, 5.2 Learning.

Experiments on Picture Representation Using Regular Decomposition

Allen Klinger, and Charles R. Dyer
Computer Science Department, School of Engineering and Applied Science, University of California at Los Angeles 90024

Computer Graphics and Image Processing 5: 68-105, March 1976

A program which uses top-down recursive partitioning of picture area into successively finer quadrants and uses logarithmic search yields tree structures containing information on such key global properties as symmetry, shape, and orientation of constituent objects or patterns. Experiments were done to merge areas, using the tree to locate the objects. Quantitative results were obtained from alphabetic letters, blocks in a scene, and the simple structures of a polyhedron. Segmentation errors caused by regular decomposition can be overcome in all cases by a relatively simple algorithm which uses the same type of concept. The resulting picture representation enables rapid access of image data without regard to position, and efficient storage.

COMPUTATION: PICTORIAL SYSTEMS

Picture Processing: 1975

Azriel Rosenfeld
Computer Science Center, University of Maryland, College Park 20742

Computer Graphics and Image Processing 5: 215-237, June 1976

354 items arranged under the following headings: Introduction; Transforms and Filtering; Compression; Enchancement, Restoration, and Reconstruction; Implementation; Pictorial Pattern Recognition; Matching and Local Feature Detection; Segmentation; Shape and Texture; Scene Analysis; Formal Models. A brief prefatory article serves as a guide to the bibliography.

COMPUTATION: PICTORIAL SYSTEMS

Computational Techniques in Visual Systems Part II. Segmenting Static Scenes

Edward M. Riseman, and Michael A. Arbib
Department of Computer and Information Science, University of Massachusetts, Amherst 01002

COINS Technical Report 76-11

Part I was concerned with the interaction between high- and Low-level systems. Part II is concerned with specific tasks of the low-level systems - feature extraction and segmentation, as well as their competition and cooperation. The emphasis is on integrated system design, with interaction of multiple processes resolving ambiguous and noisy data. A survey of processes which operate on a single static, but colored, image show how segmentation can proceed via boundary formation, and by formation of regions on the basis of color and texture cues. Experimental data are given.

Computational Techniques in Visual Systems Part I. The Overall Design

Michael A. Arbib, and Edward M. Riseman
Department of Computer and Information Science, University of Massachusetts, Amherst 01002

COINS Technical Report 76-10

Our overall goal is to define computational techniques to be used by a system in making a visual scan of a dynamic environment with which it is to interact. Here, we discuss both brain mechanisms in the visual systems of animals and humans and computer techniques for the analysis of color photographs of natural scenes. We present schemas as a formalization of the system's 'knowledge units'. This notion is helpful for our work in both the BT (Brain Theory) and AI (Artificial Intelligence) approaches. We further present specific studies -- from our own group and from elsewhere -- of subsystems of both animal and computer visual systems. We shall examine the interaction of high-level processes with low-level systems, as part of a general emphasis on integrated system design.

A Progress Report on VISIONS: Representation and Control in the Construction of Visual Models

Allen R. Hanson, and Edward M. Riseman
Department of Computer and Information Science, University of Massachusetts, Amherst 01002

COINS Technical Report 76-9

The goal of VISIONS is the segmentation and interpretation of a digitized color image of natural outdoor scenes. Multi-level data structures are used for representing both a visual model and the semantic data base of stored knowledge about the world. A flexible modular strategy controls the operation of processes which embody diverse forms of knowledge, and allows both data-directed and knowledge-directed model building. A model search space is used to store a sketch of the processing history during model formation so that limited, directed backtracking will be facilitated. A symbolic data structure (RSE for *R*egions, line *S*egmentations and *E*ndpoints) interfaces the results of low-level segmentation processes with the interpretation processes which form hypotheses about surfaces, objects and frames of visually familiar situations. The RSE structure represents syntactic two-dimensional image information while the three higher levels of representation organize semantic concepts in three-dimensional space. Utilization of the RSE structure decomposes the development of the low-level and high-level systems; it provides a clear statement of the requirements imposed on the low-level segmentation processes, and delineates the form of the data which will be the input to the high-level processes.

Precision Weighting - An Effective Automatic Indexing Method

C. T. Yu
Department of Computing Science, University of Alberta, Edmonton, Canada

G. Salton
Department of Computer Science, Cornell University, Ithaca, NY 14853

Journal of the Association for Computing Machinery 23: 76-88, January 1976

There are two principal drawbacks to current work in automatic indexing: 1) the semantic role of individual terms or concepts in query or document texts is given up in favor of formal characteristics, such as their frequency distributions, or their location on the body of a text; 2) the measurement of retrieval effectiveness is entirely experimental, with no attempt to provide mathematical proofs of the relative merits of different methods. A method of precision weighting has been developed which partially meets these objections. A precision weight is attached to each query term according to whether the term occurs primarily in documents identified as relevant to a given user query or whether it occurs in the nonrelevant documents. Given such a precision weighting system and an assumption concerning the distribution of the vocabulary accross the documents of a collection, formal proofs establish that at *every level of recall* the precision weighting system may be expected to be superior to a system in which terms in the query and document vectors are unweighted.

DOCUMENTATION: RETRIEVAL

Some Thoughts on an Interactive Information Retrieval System

D. J. H. Brown
Applied Mathematics Department, University of the Witwatersrand, Johannesburg

Proceedings of the Symposium on Information Retrieval,
University of Natal, 1975

The interaction between the subject areas of Artificial Intelligence and information Retrieval appears currently to exist only at the level of 'fact retrieval' and not 'document retrieval' Perhaps it might be possible for a natural language understanding program linked to a document retrieval system to perform selected searches for material in which the degree of sophistication, orientation etc., of the user making an enquiry can be inferred from a dialogue and used to direct searches. Some suggestons are put forward which utilize some of the concepts embodied in CONNIVER.

Development of an Integrated Energy Vocabulary and the Possibilities for On-Line Subject Switching

R. T. Niehoff
Battelle Columbus Laboratories, Columbus, Ohio 43201

Journal of the American Society for Information Science 27(1): 3-17
January-February 1976

An integrated vocabulary of energy terminology containing 24,000 entries was developed from 11 different vocabularies from government and non-government sources. All energy terms were processed with thesaurus generator software which creates reciprocal cross references and complete generics as required. It was concluded that vocabulary conversion, the ability, to retrieve all documents on a given subject from all available (and appropriate) data bases with a single query, is the best method for achieving intersystem compatibility. A prototype conversion guide (synonym table) was constructed for further study where it was found that, without any additional intellectual efforts, conversion can be increased from 28% (exact match) to 56% using exact match plus singular-plural equivalencies, plus synonym expansion.

TRANSLATION

Development, State of the Art and Prospects in Automatic Translation (*soucasny stav a vyhledy automatizace prekladu*) (In Czech)

Pitha P. Vyvoj

Synopsis 4(4): 50-81, 1973

The following issues are discussed: 1) Vocabulary and phraseological analysis. 2) Morphological analysis. 3) Syntactic analysis. 4) Semantic analysis. 5) Synthesis is a task in its own right and not just the reverse of analysis. 6) The problems of implementation of automatic translation: 7) A Brief history of automatic translation. 8) The state of the art in Czechoslovokia (essential use of a profound theoretical model, the multilevel functional model proposed by Sgall; bringing together theoretical and applied aspects). 9) Algebraic linguistics and formal linguistic models in Czechoslovakia. 10) Three conclusions: (a) practical MT is feasible but its quality and economic efficiency depend on the quality of linguistic description of the languages concerned; much theoretical work is still needed; (b) automation of any kind of information processing task which involves processing of texts requires the same means as MT; therefore work in MT should underlie this work; and (c) algebraic linguistics plays a major part in building MT systems.

SOCIAL-BEHAVIORAL SCIENCE: PSYCHOLOGY

A Simultation Model of Psychological Epistemology: Meta, Empo, and Ratio

Gregory P. Kearsley
Center for Advanced Study in Theoretical Psychology, University of Alberta, Edmonton, T6G 2E1 Canada

Behavioral Science 21: 128-133, March 1976

3 different epistemic styles are considered. *Rationalism*: does it (a piece of new information) contradict anything already know? *Empiricism*: does it agree exactly with something else? *Metaphorism*: does it contain something similar to anything known? The epistemological hierarchy of an individual involves all three epistemic styles, ordered in one of 6 possible ways. The KNOWING program (in LISP 1.5) is capable of rearranging lists of atoms, lists of lists, or lists of lists according to the different epistemic rules. , EMPO(x), METO(x), and RATIO(x) rearrange their arguments according to the three epistemic types. The function KNOWING (style 1, style 2, style 3, x) applies epistemic style 3 to list x, then style 2 to the result, and finally style 1 to that result. The function PROFILE(x) outputs the six different epistemological hierarchies of KNOWING for x. The program was tested against 8 human subjects whose epistemic hierarchies had been determined by psychological testing. The task was to rearrange 5 sets of nonsense sentences in some order which made sense. The experiment revealed the existence of major inadequacies in the current program as a psychological model.

SOCIAL-BEHAVIORAL SCIENCE: PSYCHOLOGY: PSYCHOLINGUISTICS

An Application of Concept Formation Techniques to the Learning of a Linear Language

D. J. H. Brown
Applied Mathematics Department, University of the Witwatersrand, Johannesburg

In R. Trapple and F. de. Hanika, eds., Progress in Cybernetics and Systems Research Vol. II, Halsted Press, 1975
ISBN 0-470-88476-2

In order to understand language and how it is used it is necessary to discover how it is learned. A new-born child (it is assumed) does not have at his diaposal a built-in syntax, but he does have mechanisms for collecting and assimilating information. The first stage of learning is assumed to be the building of a set of associations between physical entities and their linguistic representations. In children, there are additional problems regarding the learning of the various physical manipulations necessary to actually pronounce the appropriate word(s); this point is not pursued here. Learning of associations appears to involve two main precepts: perception of a physical entity, and identification of the context in which the entity appears. From these, the learning to be achieved involves the formation of concepts about the entities perceived. These concepts include descriptions of classes of entities and associations with contexts and linguistic representations.

The London Stage 1800-1900: A Data Base for a Calendar of Performances on the Nineteenth-Century London Stage

Joseph Donohue
Department of English, University of Massachusetts, Amherst

Computers and the Humanities 9: 19-186, July 1975

A data bank containing information about daily performances on the London Stage in the 19th Century may well be as many as a billion characters representing as many as a million performances. Each category of information entry into the computer is tagged with a unique letter, a *box code*. Information assigned to these codes includes: 1) basic information about the pieces performed on a particular data at a particular theatre or music hall, 2) supplementary general information and information on incidental performance, 3) annual or seasonal summary and commentary separate from daily calendar entries. Each record (data for a given theatre on a given day or night) begins with a Header containing date, performance, time, theatre name, of contributing editor responsible for compiling the record. Since it is desirable to be able to retrieve everything from the data base a system of indexes and subindexes has been devised, with theatres serving as Master Referents.

HUMANITIES: ANALYSIS

A Program Proposal for the Linguistic Analysis of Legal Texts in Old Spanish

Robert A. MacDonald
University fo Richmond

SIGLASH Newsletter 9 No. 1-2: 7-18, Dec. 75-Mar. 76

A discussion of the procedures for computer storage of word-byword analysis of the Alfonsine *Especulo* to be used for studies in phonology, orthography, morphology, syntax, lexicology, and, to a lesser extent, styllistics. Many examples of entry formats.

A Computer system for Conversational Elicitation of Decision Structures

Antonio Leal, and Judea Pearl
4731E Boelter Hall, University of California, Los Angeles 90024

UCLA-REP-7665, June 1976

An interactive computer program has been disigned and implemented that elicits a decision tree from a user in an English-like conversational mode. It emulates a decision analyst who guides a decision maker in structuring and organizing his knowledge about a particular problem domain. The approach centers on the realization that the process of conducting an elicitation dialogue is structurally identical to conducting a heuristic search on game trees. The program requires the decision maker to provide provisional values at each intermediate stage in the tree construction that estimates the promise of future opportunities open to him from that stage. These provisional values then serve a role identical to a heuristic evaluation function in selecting the next node (scenario) to be explored in more detail. This permits real-time rollback and sensitivity analysis as the tree is being formulated, in order to concentrate the effort on those parts of the tree which are crucial for the resolution of the solution plan.

Processes in Acquiring Knowledge

Allan Collins
Bolt Beranek and Newman, Inc., Cambridge, Massachusetts 02138

The objective of this paper is to develop a theory of Socratic tutoring in the form of pattern-action (or production) rules for a computer program. These pattern action rules are being programmed on a computer system for tutoring causal knowledge and reasoning. The paper includes twenty-three production rules derived from the data analyzed, together with segments of the data showing the actual application of the rules in different tutorial dialogues. The strategies themselves teach students: (1) information about different cases, (2) the cuasal dependencies that underlie these cases, and (3) a variety of reasoning skills. These include such abilities as forming hypotheses, testing hypoteses, distinguishing between necessary and sufficient conditions, making uncertain predictions, determining the reliability or limitation of these predictions, and asking the right questions when there is not enough information to make a prediction.

From Understanding Computation to Understanding Neural Circuitry

David Marr, and Tomaso Poggio
MIT AI Laboratory, Cambridge, Massachusetts 02139

Memo 357, $1.30, 22p., May 1976

The central nervous system needs to be understood at four nearly independent levels of description: 1) that at which the nature of computation is expressed; 2) that at which the algorithms that implement a computation are characterized; 3) that at which an algorithm is committed to particular mechanisms; and 4) that at which the mechanisms are realized in hardware. In general, the nature of a computation is determined by the problem to be solved, the mechanisms that are used depend upon the available hardware, and the particular algorithms chosen depend on the problem and on the available mechanisms. Examples are given of theories at each level.

ROBOTICS

A Logical Theory of Robot Problem Solving

Olga Stepankova
Institute of Computation Techniques, Technical Unviersity of Prague, Czechoslovakia

Artificial Intelligence 7: 129-161, Summer 1976

The concept of an image space, motivated by the SRI robot planning system STRIPS, is introduced as a formal logical counterpart to the state-space of problem solving in robotics. The main results are two correspondence theorems establishing a relationship between solutions of problems formalized in the image space and formal proofs of certain formulas in the associated situation calculus. The concept of a solution, as used in the second correspondence theorem, has a rather general form allowing for conditional branching. Besides giving deeper insight into the logic of problem solving the results suggest a possibility of using the advantages of image-space representation in the situation calculus and conversely. The image space approach is further extended to cope with the frame problem in a way similar to STRIPS. Any STRIPS problem domain can be associated with an appropriate image space with frames of the same solving power.

Robot Computer Solving System

E., William Merriam
Bolt Beranek and Newman, Inc., Cambridge; Massachusetts 02138

BBN Report Nos. 3234 and 3253

These two reports amplify and extend the work discussed in BBN Report No. 3108. Report 3234 includes several possible conceptual interpretations of the World Simulation used in the BBN robot system, as well as a description of a visual tracking algorithm which has been implemented using that simulation. Report 3253 discusses how slop and uncertainty are implemented in the World Simulation; how slop gives rise to specification error problems which must be dealt with by the robot's Cognitive system; how Command Translation takes place from the Cognitive System to the Sensori-motor System how time is dealt with; and the necessity for a Sensory Memory Management Module. Additionally, some thoughts are presented concerning the "Going to a Plcae" problem (which is essentially generalized path-planning) and the next research steps to be undertaken.

ROBOTICS

An Experimental Robot Computer Problem Solving System

E. William Merriam
Bolt Beranek and Newman, Inc., Cambridge, Massachusetts 02138

BBN Report No. 3108

The Experimental Robot computer Problem Solving system is directed toward developing a computer model of sensori-motor intelligence. As presently envisioned the system willl include a simulated world (a Martian landscape), a Sensori-Motor System, a Cognitive System, and a Human Monitor and Control System. The robot includes a motion system, a tactile sensor, and a highly-developed vision system. The simulation provides a complex animal-like organism to be used in developing the cognitive system. Detailed descriptions are given outlining the algorithms used to calculate the visual occlusion of one object by another and the obstruction of the robot's path. The completed portions of the Human Monitor and Control System are described, and a description is included of the expected next steps in the research.

PATCOSY - A Database System for the National Health Service

J. M. Kerridge
Sheffield Polytechnic, Pond Street, Sheffield S1 1WB, Great Britain

The Computer Journal 19: 98-103, May 1976

A PATCOSY (PATIENT COMPUTER·SYSTEM) database is defined by a definition program and the data is then manipulated by application programs. A definition program is created from a series of PATCOSY Definition Language (PDL) commands. The output from the program is a sequence of data maps, one for each record type in the database. The data maps are accessed by the PATCOSY Manipulation Language (PML) commands so that particular record occurences can be made available to the applications program which originated the command. Hence PML is a host language access mechanism. PATCOSY has been designed specifically for the National Health Service (GB) and therefore reflects the structure and inter-relationships which exist between the information created by the different parts of the health services when a total medical information system is considered. PATCOSY is described in relationship to the CODASYL recommendations.