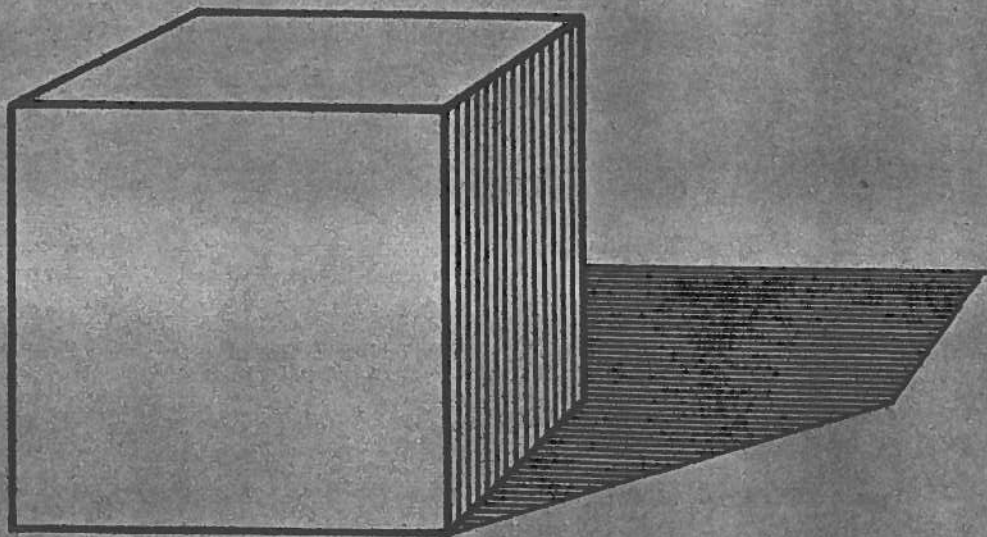


VOL. 1 / NO. 2

FEBRUARY 1975

CSCSI / SCEIO NEWSLETTER

AN OCCASIONAL PUBLICATION OF THE CANADIAN SOCIETY FOR
COMPUTATIONAL STUDIES OF INTELLIGENCE / SOCIÉTÉ CANA-
DIENNE DES ÉTUDES D'INTELLIGENCE PAR ORDINATEUR



"THERE'S STILL A GREAT DEAL
I DON'T UNDERSTAND"

ANNOUNCEMENTS

SECOND NATIONAL CONFERENCE

Canadian Society for Computational Studies of
Intelligence/Societe Canadienne pour les etudes
d' Intelligence par Ordinateur

OTTAWA May 28th - 29th

This conference has been arranged to follow a
conference on Man-machine Communications (N.R.C.,
May 26-27) in case people wish to attend both.

Sessions are being organized on the following topics. If you
wish to contribute to any of these or to present a brief paper
send an abstract (by May 1) to one of the organizers whose name
is listed beside the topic.

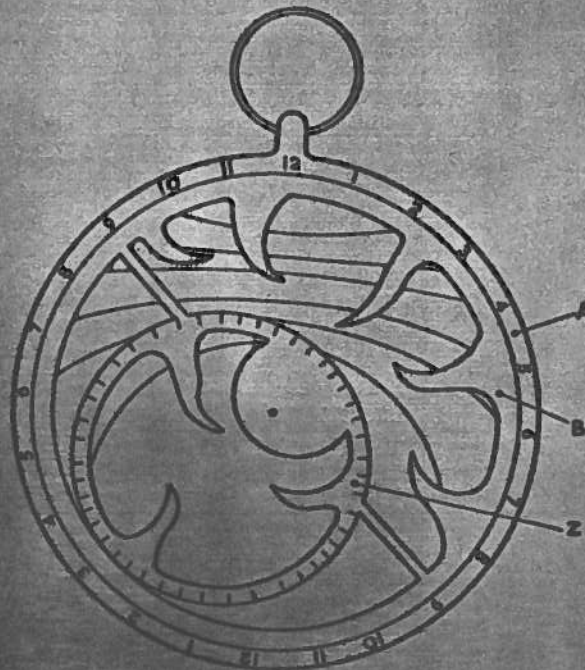
1. Cognitive modelling and computer enhancement of
education. (Zenon Pylyshy, Psychology, U.W.O., London,
Jean Gascon, Psychology, Univ. de Montreal).
2. Pattern Recognition and Picture Processing. (Tony Kasvand,
Room 345 Building M50, NRC, Ottawa), Bill Armstrong, Dept.
d'Informatique, Universite de Montreal, Wayne Davis, Dept.
Computing Sciences, University of Alberta, Edmonton).
3. Artificial Intelligence. (John Mylopoulos, Dept. of Computer
Science, University of Toronto, Doug Skuce, Dept. of Neurology
and Neurosurgery, Montreal Neurological Institute, McGill
University, Ted Elcock, Dept. of Computer Science, University
of Western Ontario, London)

C.S.C.S.I. members (as well as anyone else expressing an interest) will
be receiving further information on this conference early in May.

Hope to see you at the Tulip Festival!

TABLE OF CONTENTS

Editorial Comment	1
Letter from the (pro tem) secretary	2
Bulletin Board	3
Current Research Development	8
Artificial Intelligence Courses	17
Research Abstracts	37
Research Papers	72
List of Contributors	85



a rather ordinary Heuristic



BULLETIN BOARD

National Research Council
University of British Columbia
Université de Montréal
University of Western Ontario
Association of Literary and Linguistic Computing

LETTER FROM THE (PRO TEM) SECRETARY.

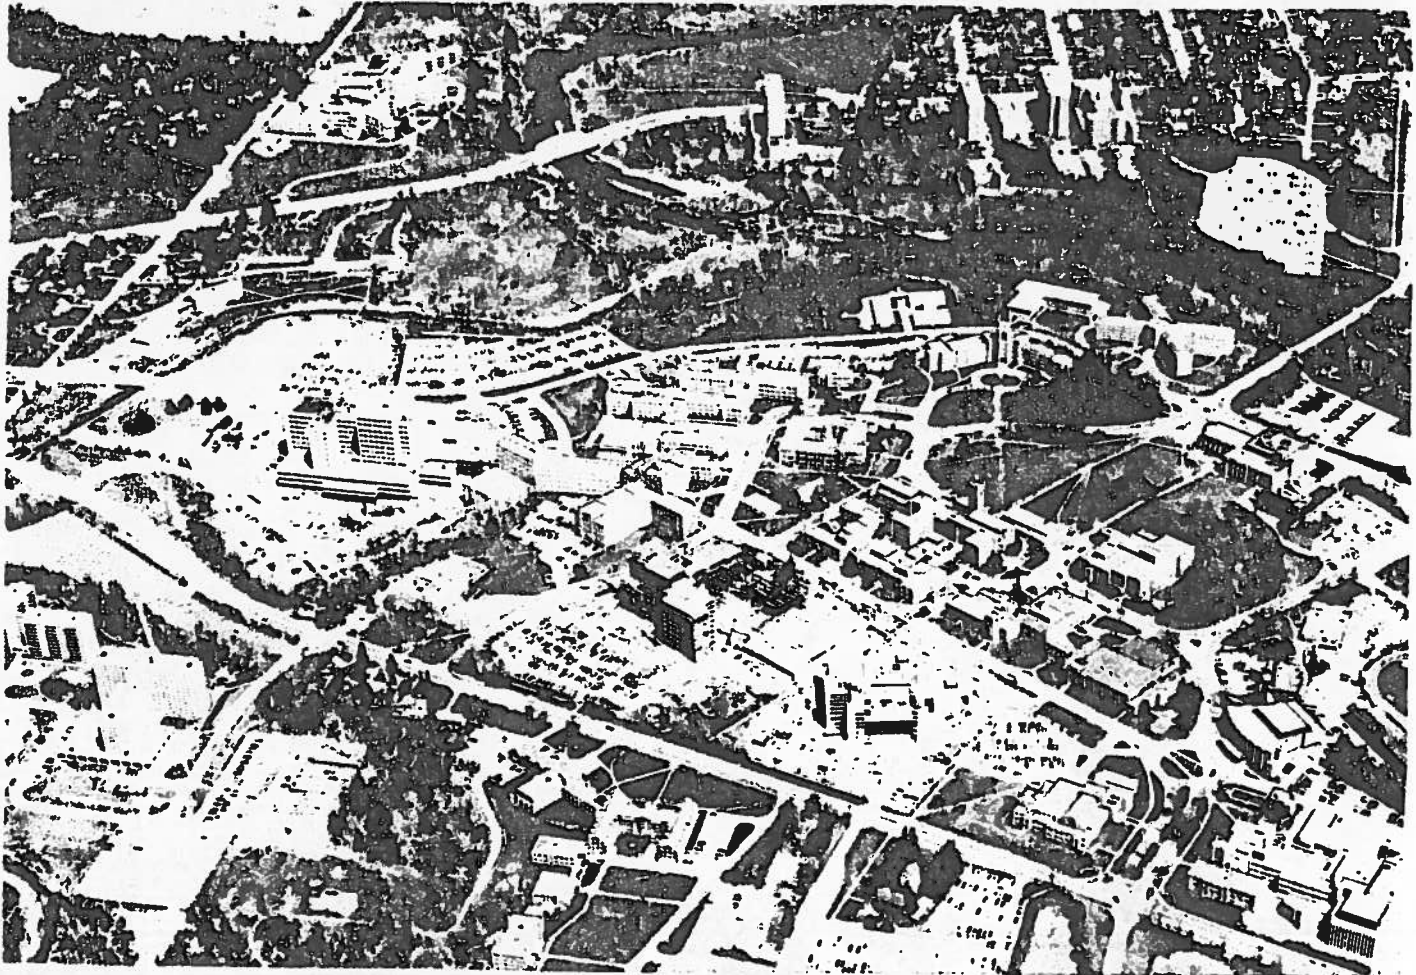
Recently I was chosen as Acting Secretary-Treasurer of our organization upon the stepping-down of John Hart. It is my duty to remind all members that their 1975 dues should be paid. Please send a cheque or money-order for \$3.00 payable to "The Treasurer, CSCSI" and mail it to

Donald Kuehner
Department of Computer Science
University of Western Ontario
London, Ontario, N6A 3K7

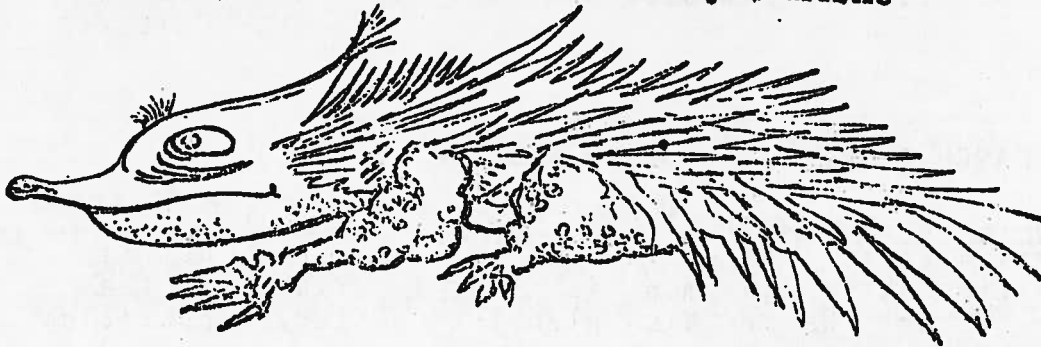
Most of this money will undoubtedly be used to defray the cost of producing this newsletter.

Any members knowing someone interested in joining C.S.C.S.I. should invite them to do so; they merely have to send us their name, address and 1975 dues.

Donald Kuehner,
Secretary-Treasurer C.S.C.S.I.



a rather ordinary Heuristic



BULLETIN BOARD

National Research Council
University of British Columbia
Université de Montréal
University of Western Ontario
Association of Literary and Linguistic Computing

This section functions as a kind of C.S.C.S.I. Bulletin board. Some of the items below were extracted from contributor's letters, others were composed specially for this section.
(the Editors)

RESEARCH & OPINIONS - Tony Kasvand (N.R.C.)

"Ongoing research, etc.: I have been studying the 3D scene segmentation as an interesting problem in itself, but also to fill the time while equipment is being interfaced to the computer system. We are also planning to study handwriting and image compression.

Opinions, etc. applicable to pattern recognition: I think that the problems normally considered are basically too simple or they are immediately simplified to bypass an obstacle. The usual explanation is, of course, that by doing simple problems we gain insight into more complex problems. My experience has indicated that the exact opposite is true, i.e. simple tricks that work on specific problems are of no use in general situations. However, I must admit that results on simple problems impress the "uninitiated" while explanations of the more general studies only produce incomprehension in the best of cases. Perhaps the "professional image" at least should be dichotomised on these lines.

Educational development: I have occasionally given a course at Ottawa University on picture processing and pattern recognition."

NEW APPOINTMENT - Alan Mackworth (U.B.C.)

"I am now ensconced as an Assistant Professor in Computer Science, teaching Computer Science 422 (Intro. to A.I.) (all year), Computer Science 502, (A.I. 1) (this term) and Computer Science 532 (Computational Vision) (next term).

I am using Introduction to Artificial Intelligence by Philip C. Jackson, Jr., Petrocelli Books, N.Y., 1974 in CpSc 422. Recommend it as text despite its idiosyncracies, typos and outright errors. It's the first textbook to cover the field properly. It'll do 'til Pat Winston publishes his MIT course notes."

RESEARCH AND DEVELOPMENT - William Armstrong (U. de Montreal)

"I am enclosing abstracts of three publications related to some A.I. work at the Universite de Montreal. There is

also an item about the formation of a Canadian company for developing a type of pattern recognition hardware."

"I hope Canadian researchers in A.I. and pattern recognition will show interest and support for our work, since it could put Canada in an advantageous position both scientifically and economically if we can develop it in this country."

"A Canadian corporation is developing special-purpose hardware for pattern-recognition applications. The system will employ adaptive tree-networks as high-speed "elementary recognizers". The basic research on tree-functions which led to this development was carried out at the Département d'informatique, Université de Montréal. Some theoretical results on tree functions have appeared as reports (please see the abstracts). Information about the hardware is not yet available to the public, however interested persons may have their names placed on the mailing list for future announcements by writing to:

DENDRONIC DECISIONS LIMITED,
1010 Ste. Catherine Street West, #640,
Montreal, Quebec.
(Attention Mr. S. Roth)."

NEW APPOINTMENT - Julian Davies (U.W.O.)

"Before coming to Western in July 1974, I was a research fellow in the Theoretical Psychology Unit, School of Artificial Intelligence, University of Edinburgh. Although interested in machine vision, most of my work at Edinburgh centred on programming languages for A.I., and on Natural Language. I designed and implemented the language POPLER 1.5, which is a PLANNER-like language suited to the POP-2-only environment that then obtained at Edinburgh. More recently, with the transfer of the Edinburgh groups to a DEC system-10 computer imminent (and now presumably accomplished) I worked on adapting the implementation of POP-2 then available for that computer to suit A.I. work (particularly POPLER 1.5). This work on POP-2 has continued at Western, and the adapted POP-2 has been renamed POP-10 (to avoid confusion).

At Western, I will be closely involved in the development of a high-level programming language for the Interdata 7/32 computer, and in maintaining POP-10. I am also interested in the problem of communications between the DEC system-10 and the Interdata, and intend to develop a 'high-level communications' package to match the high-level languages which will be used in both computers."

NEW COMPUTER/CAMERA SYSTEM (U.W.O.)

In late December, the A.I. group at U.W.O. took delivery of their new computer/camera system. A fairly complete description of this system can be found in "Equipment for a Vision Project" by E.W. Elcock in the CURRENT RESEARCH DEVELOPMENTS section of this newsletter. Lately the group, primarily through the efforts of Julian Davies, has been experimenting with the Interdata operating system and preparing the camera interface. The next step will be the development of a medium-level language on the Interdata. This language is to provide convenient and flexible programming facilities for low-level picture processing as well as interfacing well with POP-10 data structures when the Interdata/PDP-10 link is being used.

NEW (PRO TEM) C.S.C.S.I. SECRETARY

With the resignation of U.W.O.'s J.F. Hart as C.S.C.S.I. secretary, Don Kuehner has been appointed to fill that role until a new secretary is appointed at a regular C.S.C.S.I. meeting.

ASSOCIATION FOR LITERARY AND LINGUISTIC COMPUTING

We have received the good wishes of this organization for the success of C.S.C.S.I. As well as being kind enough to send us a bibliography of Polish A.I. papers, they have invited C.S.C.S.I. members to join the A.L.L.C.:

"ASSOCIATION FOR LITERARY AND LINGUISTIC COMPUTING

Chairman: PROFESSOR R.A. WISBEY

Department of German, King's College, University of London
(Founder and former Director, University of Cambridge
Literary and Linguistic Computing Centre)

This international association has been founded to further research in the field of literary and linguistic computing. It provides a means of communication for those engaged in the processing of natural languages and above all of literary texts.

REPRESENTATIVES:

Belgium	Professor L. Delatte.....LASLA, Liège
Canada	Dr. D.R. Tallentire.....University of British Columbia
Finland	Professor I.T. Piirainen.... University of Jyväskylä

France	Professor B. Quemada..... Sorbonne Nouvelle, Paris
Germany (West)	Professor W. Lenders.....Universität Bonn
Hungary	Professor F. Papp.....University L. Kossuth
Italy	Professor A. Zampolli.....CNUCE, Pisa
Netherlands	Dr. F. de Tollenaere.....INL, Leiden
Norway	Professor K. Heggstad.....Universitetet i Bergen
Sweden	Professor S. Allen....University of Göteborg
USA (East Coast)	Professor J. Raben.....Queens College, CUNY
USA (West Coast)	Dr. R. Hirschmann.....University of Southern California

**SPECIALIST GROUPS - REGIONAL BRANCHES - SYMPOSIA - SUMMER
SCHOOLS - SEMINARS**

ALIC BULLETIN: contains guest editorials, articles, reports,
correspondence, book reviews, software reviews,
bibliography, general sections.

SUBSCRIPTIONS:

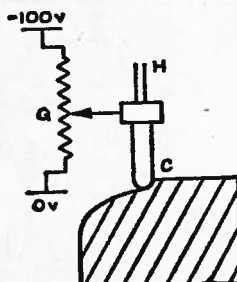
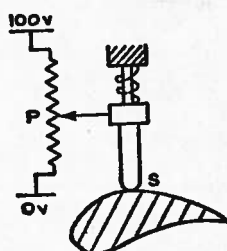
Individual Membership (includes receipt of the ALLC
Bulletin); reduced rates when attending Conferences
and Summer Schools) £3.00 for 1974 (US \$9.50)

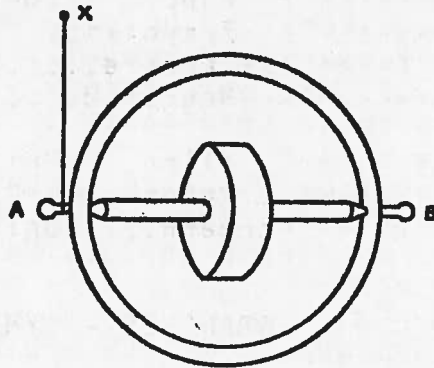
Institutional Membership (includes receipt of two copies
of the ALLC Bulletin; reduced rates when attending
Conferences and Summer Schools) £7.50 for 1974
(US \$23.50)

ALLC Bulletin only (published three times a year)
£3.00 for 1974 (US \$9.50)

For further information and an application form contact the
Honorary Secretary:

Mrs. J.M. Smith
6 Sevenoaks Avenue
Heaton Moor
Stockport
Cheshire SK4 4AW"





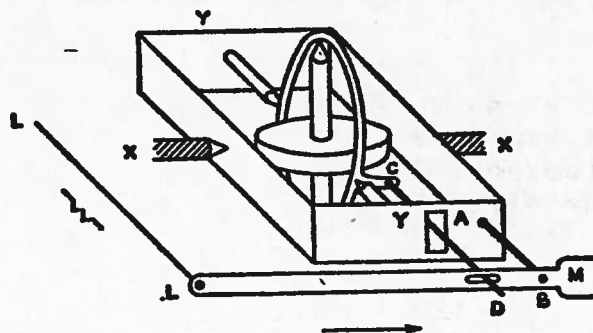
CURRENT RESEARCH DEVELOPMENTS

National Research Council

University of Guelph

University of Manitoba

University of Western Ontario



Equipment for A.I. Research at the N.R.C.
T. Kasvand
National Research Council

The N.R.C. Data Systems Division is using the following equipment in picture-processing research:
Computer hardware: Modcomp III computer, 32K core, 16 bits/word, 10^6 word disc, 2 nine track 800 bpi tapes, paper tape reader, teletype, card reader, versatec printer.
Special hardware; some being interfaced: Tektronix 61 memory scope, Hughes 639 scan converter with T.V. camera, flying spot scanner for 35 mm film transparencies (with film transport), special operator keyboard with lit pushbuttons, switches and potentiometers for analogue inputs.

Three Dimensional Colour Sculpture
Lawrence J. Mazlack
University of Guelph

The purpose of this project is to explore the use that a computer may be put to in creating three dimensional visual art. An artist working with a computer in the creation of visual art can produce work that neither machines or men could create alone.

This is an extremely new use area for computers. Computers have been heavily used in the art of music. However, the use of computers in the visual arts has been largely limited to machines serving as an aid to cartooning. This project will attempt to produce art of significant merit.

The techniques to be employed differ radically from the most commonly used form of computer generated colour pattern development, i.e. line development by function tracing with post pattern development colouring.

Initially, the visual art is to be constructed by "growing" colours by utilizing a restricted colour theory guided, stochastic technique combined with artist supplied colour development rules. A feedback loop will allow the artist to modify according to his conceptions, what the computer initially produces.

The proposed constructions are to be rectilinear patterns of varying colours. The colour patterns are to be developed by a growth technique utilizing a precedence ordering of the matrix points to be coloured. The pattern development is to be handled by a machine learning technique which will allow the pattern development to operate within colour selection constraints which can force the modification of already developed subportions of the pattern.

The final products of this project will be:

- (1) Visual art works of significance
- (2) A new artistic tool
- (3) A greater knowledge of the elements of artistic construction.

Additional Continuing Research: Heuristic data structures
(multiple entry, multiple path)

Current Research in the Department of Computer Science
D.A. Young
University of Manitoba

There are two on-going research projects in the Department of Computer Science (University of Manitoba) that are associated with artificial intelligence. One is concerned with the design of a model of the (human) retina, the other with the study of models of language processing and cognition in the brain. In conjunction with these two projects, there are two 700-level half-courses taught, "Information Processing in Animals and Machines" and "Language and Machine Intelligence".

The aim of the retina project, begun in 1968, is to study the image-encoding logic of the neural interconnectivities of the retina, in order, (1) to establish logic networks to provide more versatile automatic pattern recognition systems and (2) to contribute to the design of a prototype retinal prosthesis. A complete laboratory facility was set up last year, comprising an animal surgery and preparation room, a recording room, and an office. The recording room's Faraday cage is equipped with a stereotaxic and micromanipulator table, digital stimulus and control systems, essential functions monitoring systems, recording equipment and an image-moving projector.

The language project, begun in 1972, also has a two-fold aim: to investigate the logic of the language processing areas of the brain (especially of Broca's and Wernicke's areas and the pulvinar), in order to establish models of their neural networks such that (1) we can obtain models and algorithms for a more 'intelligent' and versatile system (essentially one that does not require syntactic or semantic analysis routines but that does incorporate 'concepts') for machine processing of natural language statements, questions and commands, and (2) a contribution can be made to the neurological knowledge of the functions of the language areas of the brain.

While we have hitherto had to rely almost entirely upon clinical data and research papers, the laboratory is now equipped for the study of evoked-potential wave-forms generated by the language and language association areas. Also, a Digital Equipment LAB 8E computer system is being installed in December, for on-line processing of waveforms (and retinal image) data. Recording from the cat pulvinar will also commence next spring, and this should supply some useful data on what the pulvinar in man is doing as a language centre.

Due to illness, no retina lab work was conducted during the period November last year to this summer; but, for the same reasons, the (until now more theoretical) language project went ahead and is now leading in progress reports. Logic models of neural networks have been produced that are capable of representing several functions of Wernicke's and Broca's areas, particularly with regard to question receiving and answering, and to pre-motor pre-articulatory processes (Broca's area). The cascade forms of logic being used appear to be amenable to applications in the visual and motor areas also, which is significant evidence in support of the correctness of some of these models.

When sufficient and tested models have been produced, on paper, it is our intention to design and build some neural networks that will be capable of some language processing at a significantly faster rate than that of general purpose computers, using a fraction of the storage commonly required for significant natural language processing, and at a fraction of the cost in computer time.

A.I. Languages at Western
Julian Davies
University of Western Ontario

A.I. Languages at Western

The University of Western Ontario offers an interactive time-sharing computer facility on a DEC system-10 with KA-10 processor and 208K words of core store. The system can accommodate 256 telephone lines at present, and it is planned to extend that to 512 lines - though not all of those terminals could be running jobs simultaneously.

The DECsystem-10 offers a variety of programming languages, high and low. From the point of A.I., there are: an implementation of POP-2 (the language developed at Edinburgh)⁽²⁾, two implementations of LISP; and an implementation of

POPLER 1.5 - a PLANNER-like language based on POP-2.

The implementation of POP-2 is called 'POP10' because the language has been adapted slightly to suit the DECsystem-10 environment. The principal changes are:

- 1) the character set is DEC 7-bit ASCII instead of the special 6-bit POP-2 code. This permits extra flexibility to users whose terminals can handle lower-case characters, and avoids confusing translations in I/O routines.
- 2) Identifiers and "words" in the system may have up to 127 significant characters instead of just 8. This is a particular advantage to those working with Natural Language.

The POP-10 system also includes an implementation of the advanced text editor developed for POP-2 by Bob Boyer and J. Moore during their time at Edinburgh⁽²⁾. At the date of writing (12th Nov. 1974) that editor is not complete, but is expected to be available for regular use within 2 to 4 weeks. The POP-10 system is based on an implementation of POP-2 by Conversational Software Ltd., which is a proprietary product, so POP-10 can not in general be made available to other computer installations.

The LISP systems both originated at Stanford; they are the old LISP 1.6⁽³⁾ and the newer 'U.C.I.' LISP⁽⁴⁾ implemented by R. Bobrow, R. Burton and D. Lewis and based on the Inter-LISP system designed and implemented for computers with paging hardware, particularly the BBN Tenex system. U.C.I. LISP itself is implemented for DECsystem-10s without paging hardware. It offers a sharable high-segment (unlike LISP 1.6), and also those features of Inter-LISP which could be implemented efficiently on an un-paged computer. These include a very much improved Break/Trace package, and complete compatibility between interpreted and compiled code. Also included is an editor based on the Inter-LISP editor for list structures in core, but there is some controversy over whether this editor is an improvement on the previous one - some users feel that it is too complicated.

The POPLER 1.5 implementation⁽⁵⁾ was developed in Edinburgh and written in POP-2 using the 'saved-state' facility of that language. It has been adapted slightly to run in POP-10, and is available from the author, though it cannot be run without a POP-2 (or POP-10) system. POPLER 1.5 has seen some use in Edinburgh for Natural Language research, and at Western for Picture Processing.

An Interdata 7/32 computer is on order for the picture processing project with 128Kbytes of memory. It is planned

to design and implement a high-level language for the Interdata, to provide data-structuring facilities similar to those in POP-2, ALGOL 68, or PASCAL. Work has not yet commenced on this project, and no details have been decided. It is expected that the Interdata and the DECsystem-10 will be run together with the Interdata dedicated to A.I., and interacting with a program in the 10, (see article on new equipment at UWO) each computer running a high level language. Facilities are being planned whereby the two high-level programs will be able to send complete data-structures to each other. It is thought that this facility will make it much easier to integrate operations in the two processors.

References

- 1) R M Burstall, J C Collins & R J Popplestone;
Programming in POP-2. Edinburgh Univ. Press 1971
- 2) R S Boyer, J S Moore & D J M Davies: The '77' Editor.
D.C.L. Memo 62, School of Artificial Intelligence,
Edinburgh 1973
- 3) L H Quam & W-Diffie; Stanford LISP 1.6 Manual. Operating
Note 28.6 Stanford A.I. Laboratory
- 4) R J Bobrow, R R Burton & D Lewis; UCI LISP Manual.
UC Irvine Information & Computer Science Technical
Report 21, 1972
- 5) D J M Davies; POPLER 1.5 Reference Manual. T.P.U.
Report 1, School of Artificial Intelligence,
Edinburgh. 1973

Equipment for a Vision Project
E.W. Elcock
University of Western Ontario

Thanks to a substantial grant of \$81,000 from the N.R.C., the Department of Computer Science has been able to order a Spatial Data Computer Eye 108 and an Interdata 7-32 which will be used as a post-processor/controller. The Interdata 7-32 in turn will be interfaced to the University Computer Center PDP-10 initially via a 4800 baud line.

The Computer Eye 108 consists of a TV camera, TV display and digitizer. Random access is available to any picture point on a 480 x 512 grid to sample its brightness level (8 bits). With the equipment ordered a complete frame could be sampled in 2 seconds. It is intended that the camera will eventually have a program controllable zoom lens, colour wheel and pan/tilt head.

The Interdata 7-32 will have 128K bytes of 750 nsec core memory and two disc drives.

The vision project has as long term goals the design and implementation of a total system which can make semantic interpretations of 3-D scenes from television camera input. The project will tackle the conceptual and implementation problems of design of the system so that it is capable of such things as:

- (a) deploying domain independent procedures which embody general models of perspective, lighting, support, etc.
- (b) accepting descriptions, in a suitable formal language, of a universe of objects and relations in a specific domain of scenes to be viewed (e.g., simple room interior). Such descriptions would encompass not only perceptual invariants of the objects, etc. but also conventional domain dependent presuppositions, e.g., that desks are often in one of a small set of preferred orientations with respect to walls and usually have a chair in front of them; that trapezoidal objects in a vertical plane are rare; that objects bear certain fuzzy scale relationships to one another; etc. These descriptions should be thought of as constituting a knowledge base to be deployed by the system in increasingly motivating low level pictorial processing as the system's partial interpretation of the viewed scene is extended.
- (c) exploiting the fact that the camera will have tilt, pan, iris, zoom and color filter controls which should allow a variety of cues to be explored in the disambiguation of depth, occlusion, surface boundaries, etc.
- (d) responding to conversational interaction via a teletype or scope in languages appropriate to both the 3-D scene and 2-D pictorial domains.

We have thought of the total functioning of such a computer vision system as being conveniently partitioned into two mutually interacting sub-systems.

1. The primary visual input is a grey level matrix of up to 480 x 560 picture points. The low level operations on this input include, besides operations such as the detection of contrast boundaries, coherent regions, etc., coping with problems of monitoring registration of successive camera scans arising from scan drift and other phenomena inherent in the output of the camera system.
2. The higher level functioning of the total system, concerned with organizing low level pictorial information into a consistent 3-D interpretation and, in this process, directing the verification by the low level system

of implications of the interpretation, etc. can only be embodied in large computationally complex programs. It is inconceivable that these programs could be written in any other than an appropriately designed very high level language. It is not necessary nor desirable that these two functions be implemented on the same physical machine. The second function demands a machine powerful enough to sustain compilers for and complex programs written in the chosen very high level language. At Western, this machine would be the University Computer Center PDP-10 and the language, already being used in pilot studies by the vision group, POPLER 1.5*. The first function demands a small dedicated machine capable of performance in camera real time, with implications for fast arithmetic, high data rate between camera and first level store and with adequate total configuration to allow a suitable operating system. Because the high level functioning operates with partial descriptions in an intermediate pictorial language rather than the primitive grey level data, the data rate between the two machines can be modest.

We consider the Interdata PDP-10 linkage, though not ideal, an adequate vehicle for sustaining this split functioning concept.

It is expected that the first year will be spent in completing the initial design of protocol and communication software for the PDP-10 and Interdata 7-32 camera controller/post processor, together with the extension of the sub-system of low level pictorial processing routines for the Interdata 7-32.

We also intend in this first year to build on our work done under a Department of Communications contract and initiate detailed studies of models of perspective and texture to be tested using the first system implementation above within the domain of simple room interiors.

* Popler 1.5 is a language which implements most of the ideas of Hewitt's "Planner" 1970, a Sussman & McDermott's "Conner" 1971. An implementation written for Western's PDP-10 was commissioned from its designer J. Davies of the School of Artificial Intelligence, Edinburgh University, under D.O.C. Grant in 1972.

Thesis Proposal: "An Artificial Intelligence Approach to
Reading Comprehension Using the Cloze Procedure"

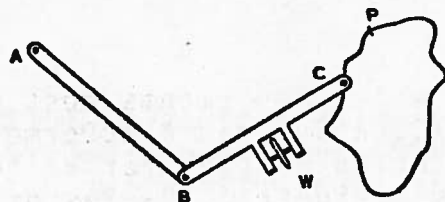
Linda L. Zimmerman
(E.W. Elcock, Advisor)
University of Western Ontario

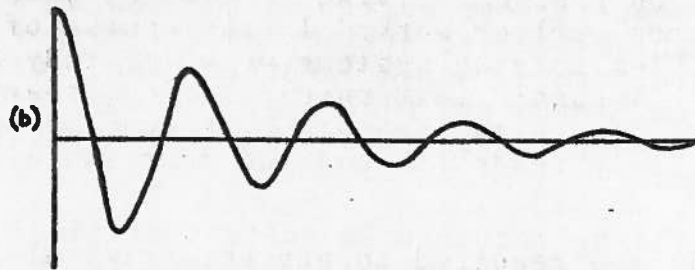
The Cloze procedure for language comprehension involves the essence of the problem of dealing with missing or ambiguous information in any context. It is proposed to explore what kind of syntactic, semantic and conceptual mechanisms must be employed in order to deal with incomplete linguistic information. Particular attention will be devoted to the disambiguation of words with multiple syntactic categories and meanings and the role of "background knowledge" in the disambiguation procedure. Also of interest will be the kind of context-sensitive mechanisms needed to discern syntactic and semantic clues from the surrounding text and relate them to the data structures representing the "content" of what has been read.

It is planned to implement such mechanisms in POP-2 and POPLER 1.5 available on the DECsys10 using a limited domain of discourse with such vocabulary and background knowledge as is necessary for basic comprehension in the problem domain. A simplified grammar will be employed, with the input text paraphrased, where necessary, to conform to this grammar. Using this tool, the investigation into linguistic applications of closure and context-sensitive mechanisms may be of some value in other pattern recognition, problem-solving environments.

The major sources for this research are:

Taylor	- Cloze Procedure
Schank	- Concept Formation
Wilks	- Preference Semantics
Winograd	- Procedural Representation of Lexical Entries
Palme	- Data Base Network Representation
Fillmore	- Case Grammar
Simmons	- Semantic Nets
Jacobson	- Connected Related Sentences





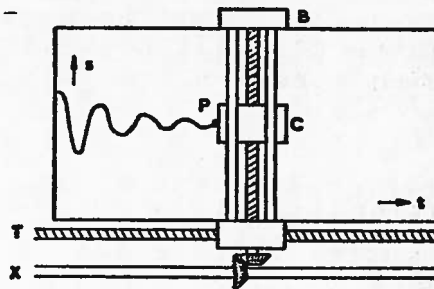
ARTIFICIAL INTELLIGENCE COURSES

University of British Columbia

University of Guelph

University of Toronto

University of Western Ontario



Artificial Intelligence I
A.K. Mackworth
U.B.C. Computer Science 502

As an introduction to Artificial Intelligence, this course will concentrate on problem-solving (including game-playing) with a bias towards earlier work. In the course of examining a variety of problem-solving systems we will study the evolution of the main issues: task specification, search, problem reduction, advice, learning, planning, analogy, generality/competence/efficiency tradeoffs and the representation and use of knowledge.

Students will be required to present critical reviews of important papers and to carry out a project to be implemented, probably in LISP.

Text: Nilsson, N.J. (1971) Problem-solving Methods in Artificial Intelligence. McGraw-Hill.

References:

Feigenbaum, E.A. and Feldman, J. (eds) (1963) Computers and Thought. McGraw-Hill.

Minsky, M. (ed.) (1968) Semantic Information Processing. MIT Press.

Slagle, J.R. (1971) Artificial Intelligence: The Heuristic Programming Approach. McGraw-Hill.

Prerequisite:

Sufficient computing experience (e.g. CpSc 310) or the consent of the instructor.

Corequisite:

CpSc 509. This requirement will be waived for students with sufficient programming experience in LISP.

Computational Vision
A.K. Mackworth
U.B.C. Computer Science 532

Computational methods for picture interpretation comprise the focus of this course. Such theories are traced from the early visual pattern recognition systems through the invention of picture description languages to recent work in scene

analysis. The work described ranges from theoretical result to commercial device; the scene domains include printed and handwritten text, formatted arithmetic expressions, nuclear particle tracks, chromosomes, human faces, road scenes, the surface of Mars and the blocks world of robotics. Peripherally related areas from other areas of artificial intelligence, conventional picture processing, graphics and mammalian vision systems are explored where relevant.

Text: None

References:

Duda, R.O. and Hart, P.E. (1973) Pattern Classification and Scene Analysis Wiley.

Grasselli, A. (ed.) (1969) Automatic Interpretation and Classification of Images Academic Press.

Kaneff, S. (ed.) (1970) Picture Language Machines Academic Press.

Rosenfeld, A. (1969) Picture Processing by Computer Academic Press.

Uhr, L. (ed.) (1966) Pattern Recognition Wiley.

Winston, P.H. (ed.) (1973) Progress in Vision and Robotics A.I.T.R. - 281, MIT.

Prerequisites: Sufficient computing experience (e.g. CpSc 310) or consent of the instructor.

Artificial Intelligence
Lawrence J. Mazlack
University of Guelph Computer Science 27-450

There are two courses at Guelph that contain elements of artificial intelligence. One, 27-450, is entirely devoted to AI, and is therefore entitled, "Artificial Intelligence". The other, 27-443, deals with advanced data structures and spends an appreciable amount of time dealing with data structures for AI type problems. Both are taught solely by the author. This paper will deal only with 27-450 (AI).

The 27-450 course examines some of the special topics in the areas of artificial intelligence. Exactly what is meant by "artificial intelligence" is the subject of an extended classroom definition. However, broadly speaking, artificial intelligence is defined as the active property that would enable an artificial entity (i.e., an entity not directly

derived from a life process) to exhibit activities which are said to be the product of intelligent action.

Several attributes are considered to be connected with an entity capable of intelligent action. They are: communication in a natural language, the capacity to deal with external objects (recognize, respond), the ability to draw inferences from what it sees, the potential to know and remember, the capability to solve problems, and an active intention to maximize its values. This course concentrates on two basic elements. They are (1) the learning process and (2) man-machine communication. Additional and sometimes related topics such as brain modelling, decision processes and interactive processes are examined.

The principal course texts are Feigenbaum and Feldman, Computers and Thought and Slagle, Artificial Intelligence. Nilsson, Artificial Intelligence and Uhr, Pattern Recognition are also heavily utilized. Required reading is assigned from various journals and the following:

Jacque, Computer Diagnosis
Kanal, Pattern Recognition
Mendel and Fu, Adaptive, Learning and Pattern Recognition systems
Minsky, Computation: Finite and Infinite Machines
Newell, et. al., Speech Understanding Systems
Spencer, Game Playing with Machines
Tsyppkin, "Adaptation and Learning in Automatic Systems"
Proceedings of the Second International Joint Conference on Artificial Intelligence
Proceedings of the Third International Joint Conference on Artificial Intelligence

Learning, pattern recognition and game playing are considered to be the most important aspect of the course. To motivate student learning in this area, 50% of the student's final grade is apportioned to student performance in a game that is at least partially a pattern recognition problem, namely the ancient Chinese/Japanese game of Go-Moku. The largest portion of the student's evaluation derives from how effective his program is against other students' game programs.

The course evaluation is:

A learning, game playing program	- 50%
Short paper	- 5%
Turing Machine programs	- 3%
Midterm	- 10%
Two special topic presentations	- 12%
Final exam	- 20%

The short paper has the topic: What is required for the thought process. It is not to be addressed to the question: "Can a machine think?". Any suitable aspect of the subject may be discussed from any viewpoint.

The Game Program

The game that is to be played is the ancient game of GO-MOKU as described in Lasker, Go and Go-Moku, Dover, 1960. This game can be considered an expanded version of Tic Tac Toe.

Each student writes an interactive Go program.

Each program must be capable of starting first or second.

Game Program Rules

The following rules apply:

- (1) The normal rules of Go-Moku will apply, except that play will take place on a 13 x 13 board.
- (2) In order for a program to win, it must recognize and declare that it has a winning position.
- (3) The APL program must announce its move in the form: "THE JOHN SMITH PROGRAM MOVES TO A-11" where A-11 signifies a given coordinate. The coordinates are alphabetic horizontally (A through M) and numeric vertically (1 through 11). The program must accept counter moves by the simple input of another coordinate.
- (4) The evaluation of programs will be derived from the results of play between the programs utilizing humans as intermediaries.
 - (a) There are two rounds of play between the competing programs.
 - (b) Each round consists of every program playing every other program twice. Each program has the opportunity to move first.
 - (c) Each program is limited to 15 minutes of clock time. Use of more than 15 minutes results in game loss.
- (5) Illegal Moves
 - (a) If a player makes an illegal move and it is recognized as such by his opponent, the opponent wins.

- (b) If an illegal move is not recognized, play will proceed from that point. The move will be considered to be legal immediately after the second player fails to recognize its illegality.
 - (c) If an unrecognized illegal move causes one stone to be placed upon another stone, the first stone will be removed from the board.
- (6) Claiming a Win
- (a) A player does not win unless he claims a win. The first player to validly claim a win for his own position, wins the game.
 - (b) If a player fails to claim a win, he may do so at a later time, unless his opponent develops a valid winning position and claims a win.
 - (c) If a player fails to claim a win, his opponent may
 - i. Proceed with his own game, ignoring his opponent's situation.
 - ii. Point out the winning situation and claim a draw.
 - (d) Anyone falsely claiming a winning position for anyone loses the game.

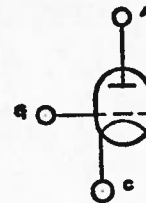
Student Game Program Evaluation

first round 17% of mark

- 5% - documentation (theory)
- 12% - result of play

second round 33% of mark

- 8% - documentation (theory)
- 25% - result of play



Result of Play Distribution:

- (a) For each round, each player is assigned 2 X (number of people in course - 1) points.
- (b) Every time a player wins, he gains one of his opponents points.
- (c) Every time a player loses, he loses one of his points to his opponent.

- (d) No point exchange for draws.
- (e) Player's mark contribution for "result of play" from a round is:

$$\left\{ \begin{array}{c} 12 \\ \text{or} \\ 25 \end{array} \right\} \times \left(\frac{\text{player's points at end of a round}}{2 \times (\text{number of people in course} - 1)} \right)$$

A Thinking and Reasoning Course in a Psychology Program
Peter A. Reich
University of Toronto Psychology 222F

As a psychology course, the course on Thinking and Reasoning at the University of Toronto emphasizes psychological research on the thinking process, while also covering material on how to think. The organizing principle of the course is to start with small circumscribed problems and to work toward bigger and more complex problems. The sequence of topics thus starts with syllogisms; it moves to school-type story problems and recreational mathematical puzzles, and then to illdefined problems. From these problems, which all concern one person manipulating inanimate objects and concepts, it moves to two-person games and their relation to international situations and wars, and from there to n-person games and problems of legislative voting. Within each topic there are two parts: the first part deals with how to do it as described by the experts: the logicians on syllogisms, mathematicians such as Polya on problems, artificial intelligence computer science people interested in ill-defined problems and game playing, and game theorists on 2-person and n-person games. The second part of each section deals with psychological research on how people really behave, as opposed to how the experts say they should behave. The course includes weekly home-work assignments which include making problem-solving procedures explicit by learning to program computers (to run a maze, for example), working anagrams and math puzzles and solving chess problems. The purpose of this is not so much to give the students practice solving problems, but to give them the opportunity to introspect on the problem-solving process.

While there has been no organized evaluation of the course, annually increasing enrollment indicates that it is successful in spite of its reputation for being a hard course. From the point of view of succeeding to teach thinking, whatever success it may have in that regard probably must be tempered by the fact that it is an optional course and thus the student population taking the course is self-selective.

Enrolment Figures:	69-70	35
	70-71	77
	71-72	82
	72-73	93
	73-74	113

PSY. 222F: Thinking and Reasoning - Fall 1973
Outline and Required Readings (492 pages)

I. Introduction to computer simulation(28)

S14F What is a computer, No readings.

S17M Homework 1: Learn to use keypunch/Tutorial:
Lesson 1 in SPITBOL.

S19W Turing. Computing machinery and intelligence. In
Feigenbaum & Feldman, Computers and Thought,
McGraw-Hill, 1963, 11-38 (28).

(Opt) Meet Shakey, the first electronic person.
Life Magazine, Nov. 20, 1970 pp 58B-68 (8).

(Opt) Weizenbaum. ELIZA - a computer program for
the study of natural language communication bet-
ween man and machine.
Communications of the ACM, 1966, 9, 36-45.

II. Deductive reasoning (99)

S21F Beardsley, Thinking Straight, Prentice-Hall, 1966,
46-78 (33).

S29M Homework 2: Syllogisms

S26W Chapman & Chapman. Atmosphere effect reexamined.
Wason & Johnson-Laird, Thinking and Reasoning
Penguin, 1968 (this reference will be referred to
by "W" from now on), 83-92 (10).

Henle - On the relation between logic and thinking.
W93-107 (15).

S28F Begg & Denny. Empirical reconciliation of atmos-
phere and conversion interpretations of syllogis-
tic reasoning errors, Journal of Experimental
Psychology, 1969, 81, 351-354 (4).

Ceraso & Provitera. Sources of error in syllogistic
reasoning. Cognitive Psychology, 1972, 2,
400-410 (11).

03M Homework 3: first computer program/tutorial Lesson 2 in SPITBOL.

010W Johnson-Laird. The three-term series problem.
&012F Cognition, 1972, 1, 57-82 (26).

III. Inductive reasoning (38)

010W Gilson & Abelson. The subjective use of inductive inference. W175-190 (16).

(Opt) Abelson & Kanouse. The subjective acceptance of verbal generalizations. In Feldman, Cognitive Consistency: Motivational Antecedents and Behavioral Consequents. Academic Press, 1966, 171-197.

(Opt) Abelson & Carol Reich. Implicational molecules: a method for extracting meaning from input sentences. In Walker & Norton, Proceedings of the International Joint Conference on Artificial Intelligence, 1969, 641-647.

012F Wason. On the Failure to eliminate hypotheses: a second look. W165-174 (10).

Smedslund. The concept of correlation in adults. In Duncan, Thinking: Current Experimental Studies. Lippincott, 1967, 432-443 (12).

015M Homework 4: Second computer program, due Review for test.

017W Test 1: Covers all material outlined above or presented in lectures.

IV. Problem solving experiments (55)

019F Maier. Reasoning in humans 11. W17-27 (11)

Dunker. On problem solving. W28-43 (16)

Birch & Rabinowich. The negative effect of previous experience on productive thinking. W44-50 (7)

022M Test 1 returned and discussed.

024W Adamson. Functional fixedness as related to problem solving. W51-55 (5).

Luchins & Luchins. New experimental attempts at preventing mechanization in problem solving. W65-80 (16)

V. Problem solving: observations and discussions (129)

026F Polya, How to Solve It. Doubleday, 1957, 1-35 (36).

Hunter. Mental Calculation. W341-352 (12).

029M Homework 5: Third computer program due, Fourth program discussed.

031W Wertheimer, The area of a parallelogram, In Productive thinking, Harper & Row, 1959, 13-78 (66).

N2F Reitman, Autobiography of a fugue. In Cognition & Thought, 166-180 (15).

(Opt) Reitman. Cognition and Thought, Wiley, 1965, 125-165.

N5M Homework 6: Polya problems.

VI. Introduction to game theory (65)

N7W Rapoport. Games. In Rapoport, Two-Person Game Theory, University of Michigan Press, 1966
(this reference will be signified by "R" from now on), 13-21 (9).

Rapoport. Utilities. R22-38 (17).

Rapoport. Strategy. R39-38 (17).

Rapoport. The game tree and the game matrix.
R47-53 (7).

N9F Rapoport. Dominating strategy and minimax.
R54-62 (9).

Rapoport. Mixed strategy. R63-77 (15).

N12M Homework 7: Game theory. Review for test

N14W Test 2

VII More on game theory (34)

N16F Rapoport. Solving the two-person zero sum game
R78-82 (5).

Rapoport. Opportunities and limitations. R186-214
(29).

N19M Test 2 returned and discussed.

N21W Non Zero-Sum games (No readings)

N26 Homework 8: Game theory.

N28W Guest lecture: Prof. Rapoport will discuss the psychology of the Prisoner's Dilemma and other games.

D5W Guest Lecture: Prof. Rapoport will discuss n-person games.

VIII Computer Models (44)

N23F Manis, Computer simulations of thought. In Manis, An Introduction to cognitive psychology, Brooks-Cole, 1971, 203-216 (12).

Newell & Simon. GPS, a program that simulates human thought. In Feigenbaum & Feldman, Computers and Thought McGraw-Hill, 1963, 279-293 (15).

N30F Shannon, A Chess-playing machine. Scientific American, Feb, 1950, 182, 48-51 (4).

Zobrist & Carlson. An advice taking chess computer. Scientific American, June 1973, 228, 93-105 (13).

(Opt) Samuel. Some studies in machine learning using the game of checkers. In Feigenbaum & Feldman, Computers and Thought McGraw-Hill, 1963, 71-108 (38).

(Opt) Greenblatt, Eastlake & Crocker. The Greenblatt chess program. Proceedings of the FJCC, 1967, 801-810 (10).

D3M Homework 9: Maze problems due, discuss chess machines.

D7F Newell & Simon. An example of human chess play in the light of chess playing programs. In N. Weiner & J.P. Schede, Progress in Biocybernetics 2, Eisevier, 1965, 19-75.

(Opt) Chase & Simon. Perception in chess. Cognitive Psychology, 1973, 4, 55-81.

Evaluation

Test 1	25%
Test 2	25%
Homework & Tutorial	15%
Final Exam	35%

Human and Machine Perception
Zenon Pylyshyn
U.W.O. Psychology 550
U.W.O. Computer Science 540
York University Psychology 675c

This course is an experimental venture running jointly among three departments. Its goal is to introduce students to the problems of perception from a computational point of view bringing in insights from psychology, neurophysiology as well as artificial intelligence. Because of the heterogeneity of backgrounds special attention has to be paid to fundamentals. We considered this venture to be worth attempting because there is great overlap in the questions being asked in the three disciplines if not in the way in which answers are sought, and frequently great ignorance of the approaches taken outside one's own field. It remains to be seen whether, given the current state of the art as well as the backgrounds of the participants, it was all worthwhile. It appears at least to be of sufficient interest that some 8 students (2U.W.O. Psych., 3 U.W.O. C.S., 3 York Psychology) and about the same number of faculty attend regularly and most of them make the 125 mile trip between York and Western every third week.

Because material relevant to the course is widely dispersed throughout various publications, theses, and unpublished manuscripts a collection of readings was made available to students. The primary readings include those listed below.

In addition to lectures the course provides an opportunity for students to present seminars on selected topics (both human and computer) and to carry out projects on the computer. Details of the latter have not been worked out at the time of writing because of an uncertainty concerning what computational resources will be available next term. We do have some digitized pictures of office scenes and geometrical drawings and a package of POP-2 (now POP10) programs to access them but the video camera and Interdata computer should replace those if delivery is taken within the next few weeks.

COURSE OUTLINE: Human and Machine Perception

1. What (if anything) is the problem?

-The problem of intimate familiarity and naturalness before we can ask meaningful questions we must locate a puzzle (as opposed to a mystery): We must find some "confrontation" (CF I. Howard) between two ways of seeing the same thing so that there is

something that demands explaining; or we must set up a specified goal which appears to be achievable by some methodology.

-Hence preoccupation with "optical illusions"

-but how to formulate questions about normal perception?

The attempt to devise a mechanism which "perceives" is one way to sharpen questions.

II. Psychological Traditions:

(1) Biological: "Explanation" occurs when neurophysiological events accompanying perception can be described in detail. (But does this answer psychological questions?)

(2) Behavioral:

-Associationism -- No "mental" events except as shorthand name for classes of stimuli or responses.

-Radical Behaviorism (Skinner) -- Atheoretical "engineering" approach concerned only with prediction and control of gross behavior.

(3) Phenomenological

-Structuralism -- "sensations" as atomic building blocks.

-Gestalt -- perception "wholistic", "laws of pregnanz" "psychological determinism"

(4) Information Processing Approach

-(1) to (3) are sources of evidence but the theory is at another "functional" level of analysis. Notions of "mechanism" (a la turing) and "symbol" are primary.

III. Some Special Conceptual Problems Plagu ing Perception

(1) The physical/symbolic squeeze

There are two distinct classes of concepts involved when we talk about perception. A major problem is to bridge these two vocabularies with a theory.

-Type 1: The Physical vocabulary

-E.G. Energy, wavelength, intensity, gradient,... this vocabulary is in the language of Physics. It reverts to measurements which can be made by using "conventional" instruments, or to "simple"

functions of such measurements (e.g. log of intensity). The measures apply to a point in some "space" - i.e. are "local" in a special Euclidean sense.

-Type 2: The Symbolic Vocabulary

-E.G. Tables, chairs, legs, top, edge, shadow, corner, in-front-of, supports, occludes, abuts,... this is the everyday vocabulary of objects, parts, relations, and other "interpreted" aspects, relations, or properties. It includes not only things for which we have words in our natural language but also cognitive concepts for which we do not have words at present - but could develop words or phrases for. This vocabulary type implies a form of "articulated" or symbolic system with each symbol standing for (or referring to) an equivalence class of possible sensory events. We shall refer to it in later discussions as a "structured description".

*The thesis of the present approach is that perception involves the construction of descriptions, using Type 2 concepts, from data which we must describe using Type 1 concepts.

Bridging the gap from the physical measures to the psychologically relevant measures is the task of the sub-discipline called "Psychophysics". Bridging the gap between the two symbolic domains in the machine vision case has been the preoccupation of a stream of workers which is usually traced from Roberts through Guzman, Clowes, Huffman, up to the elegant work of Waltz (in the readings). In the human vision case this gap has been the concern of many investigators -- most notably in the Gestalt tradition. We shall examine the views of Hochberg and of J.J. Gibson.

(2) The top-down/bottom-up squeeze (sometimes called the parsing problem)

Do we see the parts (e.g. edges and surfaces) of an object such as a table and then infer what the object is? Or do we first perceive the object and hence infer that these other parts must be edges and surfaces? The Gestalt movement came down strongly on the latter side. But there is a paradox here for how can we recognize a whole which depends on the parts without having recognized the parts?

(This question is related to the part-whole problem discussed in many disciplines. Minsky and Papert showed that the slogan "The whole is more than the sum of the

parts" or exactly what is meant by "sum". In those cases where one can specify these precisely enough (e.g. in devices called "perceptrons") there are some properties which are not deduceable from parts and others that are (see the brief discussion in the Minsky and Papert reading).)

- (3) There are a number of other "squeezes" which recur as perennial problems in perception. They are listed below but will be discussed later.
- Analog-digital
 - Parallel-serial
 - Wholistic-articulated
 - Imaginal-verbal

111. Computational approaches: "Low Level" processing

0. Transduction (machine case)
The "picture function"; digitization; pixles; T.V. input sources of technical harassment: noise.
1. Global pre-processing: picture transforms
 - Filtering, restoring, enhancement (sharpening, smoothing) (Why? seductiveness of "looking clearer")
 - Transforming into a new space
 - Gradients (first derivative)
 - Laplacian (second derivative)
 - Psychological evidence (bands)
 - Neurophysiological evidence (lateral inhibition)
 - Spatial frequency domain
 - Psychological relevance
 - Spatial filtering
 - Parameter spaces: The hough transform (rho-theta space) (Clowes' and Duda/Hart application to be covered below)

(What does it buy us?)
2. Pattern Recognition
 - The problem: Partition a pattern set and assign a symbol (name)
 - The template approach (wholism at its best)
 - obvious advantages, obvious limitations
 - psychological arguments (suggestive)
 - Hayes, Posner, Shepard, Cooper
 - implementations
 - simple binary case (noiseless)
 - grey-level case (with noise)
 - correlograms and optical methods
 - holography

- Sub-templates and related methods
(not-too productive tradition of probabilistic,
random connected sampling and learning
recognizers)
- Bledsoe and Browning
- Uhr and Vossler
- Pandemonium and Parallel Processing
- The Infamous Perceptron
- Psychological evidence for this approach
 - Hebb and stabilized retinal images
- Machine models utilizing grey-level picture
function.
 - Line finders and followers
 - Clowes line finder, Hueckel operator, Kelly
planning line finder, Pringle "accomodating"
line finder, intro to Shirai's contextual
system
 - Systems with internal models: Roberts,
Falk.
 - Other finders (e.g. region finders)

IV. Neurophysiological approaches

- A. Feature analysers and the hierarchical approach
 - (or, from bug detectors to grandmother cells)
- 1. Historical antecedents (date to mid 1950's)
 - Sutherlands "Stimulus Analysing Mechanisms"
 - Uttley's Conditional Probability Computer
- 2. Suggestive psychological evidence
 - Confusion matrices
 - Neisser's rapid search experiments
 - Fixed retinal images
 - Afterimages (especially specific motion
aftereffects)
 - Masking
- 3. Neurophysiological evidence ("The Single Cell
Revolution")
 - Single cell recording technology
 - Letvin et. al. on the frog's eye
 - Single cell receptive fields (centre-surround)
 - Hubel and Wiesel's hierarchy of cells (from
simple to complex to hypercomplex)
- 4. Related computer-based pattern recognition systems
A number of "pattern recognition" models of the
1960's directly parallel the currently fashionable
neural models. For example Barlow's discussion
sounds very much like Uttley's system as well as

the adaptive recognition schemes of Bledsoe and Browning, Uhr and Vossler or Rosenblatt's "perception" (with some similarity also to Selfridge's "pandemonium")

5. General problems with the hierarchical approach
- Neurophysiological counter-evidence: Parallel paths
 - Particular sets of features neither necessary nor sufficient for perception
 - Cf Hochberg's examples, closure figures, cognitive contours, tactile & Haptic perception
 - Some feature-like analysis necessary because of the problem of decoupling the organism from a rapidly fluctuating environment.
 - Context dependence of many "features" (or at least of extraction parameters such as thresholds, size of domains)
 - the role of attention ("noticing")-Shepard
 - making use of redundancy (Miller)
 - Limitation of "hyperspace partitioning" approach alone
 - Introduction to the importance of "Rule Governed" systems for perception
 - the analysis-by-synthesis model
 - Introduction to the importance of "Knowledge of the World" for perception
 - e.g. The speech perception and language comprehension projects.

- B. Other neurophysiological approaches (briefly noted)
- Neurophysiological isomorphism and field models (Gestalt)
 - Frequency domain models (currently fashionable in some neurophysiological circles) and its logical extrapolation to holographic memory models (see Pribram's "Languages of the Brain")
 - Distributed computation models (McCullough; Arbib). A grand idea whose time has not yet arrived.

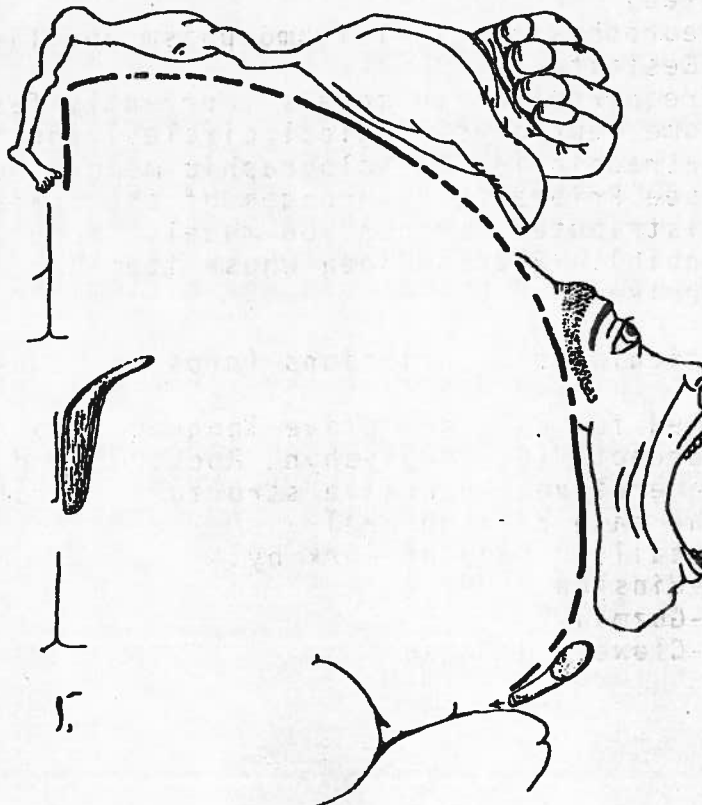
V. Articulated descriptions (pros and cons)

- Need for a "descriptive language" to represent percepts (C.F. Pylyshyn, Rock, Howard, and the Generative, recursive structures parallel to the case of language)
- Detailed study of work by:
 - Winston
 - Guzman
 - Clowes, Huffman

- Roberts, Falk
- Waltz
- (above to occupy about 1/4 of the course)
- The problem of cognitive representation
 - arguments for "analogical representations" or "mental images" as forms of non-propositional non-articulated (or Sloman's "Non-fregean" Systems)
 - Shepard's Mental Rotation Studies and Related Chronometric Studies
 - "Mental Comparison" studies and the idea of "Prototypes"
- The value and role of diagrams in problem solving
 - U.W.O. Geometer (Arbiter)
- Relation between perception and cognition
- Perception as (1) Transformation (2) Extraction (3) Problem-solving (4) Parsing (5) Construction of an Internal Structural Description (6) All of the above (7) None of the above (Choose only one).

VI. Other topics covered

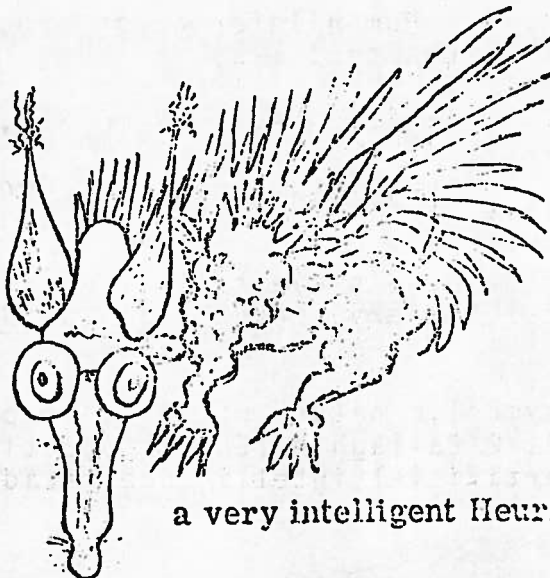
- Perception of three dimensions
 - Stereopsis; Monocular Cues (Gibson)
- Spatial Frequency Analysis (Human & Machine Cases)
- Analysis of Texture (Human & Machine Cases) -
- Speech Perception (Comparison with Vision)



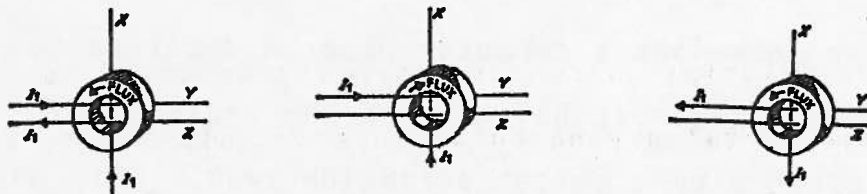
- Barlow, H.B. Single units and sensation: a neuron doctrine for perceptual psychology. Perception, 1972, 1, 371-394.
- Clowes, M.B. On seeing things. Artificial Intelligence, 1971, 2, 79-116.
- Duda, R.O. and Hart, P.E. Pattern classification and scene analysis. Wiley, 1973, Chapters 7,9,12.
- Elcock, E.W. A report on the U.W.O. vision project (unpublished).
- (s) Falk, G. Interpretation of imperfect line data as a three-dimensional scene. Artificial Intelligence, 1972, 3, 101-144.
- Gibson, J.J. The senses considered as perceptual systems. Houghton-Mifflin, 1966. (pp. 266-286).
- Guzman, A. Computer recognition of three-dimensional objects in a visual theme. MIT-TR-59, 1968.
- Hochberg, J. In the mind's eye.- In R. Haber (Ed.) Contemporary theory and research in visual perception. Holt, Rinehart, & Winston, 1968.
- Howard, I.P. Perceptual schemata and descriptive functions (mimeo, 1974).
- Huffman, D.A. Impossible objects or nonsense sentences. In Machine Intelligence 6.
- (1) Kaufman, L. Sight and Mind. Oxford University Press, 1974.
- (1) Kelly, M.D. Edge detection in pictures by computer using planning. MI 7.
- (1) Lindsay, P. & Norman, D. Human Information processing. Academic Press, 1972 (Chapters 1-3)
- Minsky, M. & Papert, S. M.I.T. A.I. Memo 252, 1972.
- (1) Neisser, U. Cognitive psychology. Appleton-Century-Crofts, 1967 (Chapter 3).
- Pylyshyn, Z.W. What the mind's eye tells the mind's brain: a critique of mental imagery. Psychological Bulletin 1973, 80, 1-24
- Pylyshyn, Z.W. The symbolic nature of mental representations. In S. Kaneff and J. O'Callaghan (Eds.) Objectives and methodologies in artificial intelligence. Academic Press, 1975.

- (s) Roberts, L.G. Machine perception of three-dimensional solids. In Tippet, et al., (Eds.) Optical and Electro-optical information processing.
- (1) Rosenfeld, Picture processing by computer. Computing surveys 1961, 1, 146-176.
- Shira, Y. A context-sensitive line finder for recognition of polyhedra. Artificial Intelligence, 1973, 4, 95-119
- Shepard, R. Studies in the form, formation and transformation of internal representations (manuscript).
- (1) Uhr, L. (Ed.) Pattern recognition. Wiley, 1966.
- Waltz, D.L. Generating semantic descriptions from drawings of scenes with shadows. M.I.T. (A.I.-TR-271), 1972.
- Winston, P. The MIT robot. In Machine Intelligence 7.
- (1) Zusne, Visual perception of form. Academic Press, 1970.

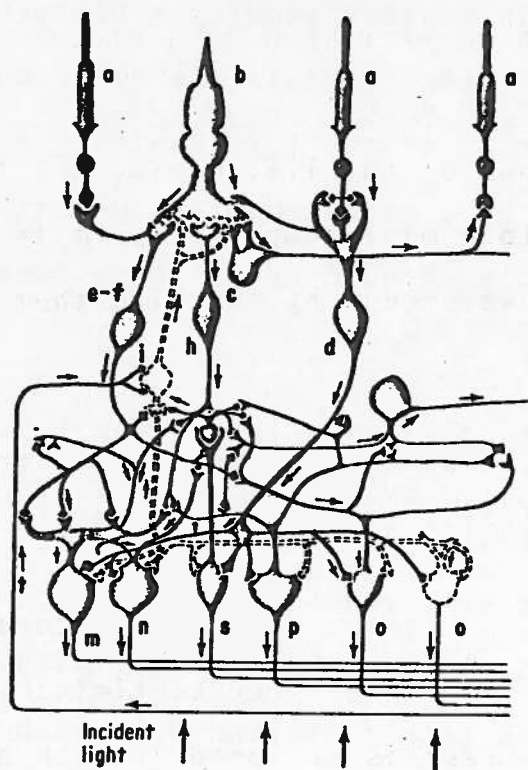
(s) -- student seminars
(1) -- source for lectures



a very intelligent Heuristic



RESEARCH PAPERS



CAN A COMPUTER REPLACE A DEVELOPMENTAL PSYCHOLOGIST?

Jean Gascon

Université de Montréal

This paper describes a computer program destined to replace a developmental psychologist in the analysis of children's protocols on a weight seriation task. This paper will not stress the programming considerations, but rather the various reasons underlying such an enterprise and its possible implications for developmental psychology.

The task is well-defined: Children, ages six to ten, come into an experimental room, one at a time, and are presented with a television monitor showing a picture of a two pan balance, together with a certain number of blocks (5) to be seriated. In front of the T.V. monitor, a touch-sensitive glass hooked up to a mini-computer, permits the child to make the blocks move around by touching them on the screen.

The instructions are the same as those used in the usual piagetian weight seriation experiment.

Such an elaborate set-up permits the exact recording of the children's actions, together with precise timing information. Secondly, the task is stripped of such extraneous

factors as kinesthetic information that often confounds the data. Finally, this apparatus permits on-line protocol analysis by connecting the mini-computer to a fast and powerful time-sharing computer.

Such a configuration has been realized and is now undergoing serious testing. The most interesting part of the program has to do with the understanding of children's protocols and the identification of the strategy best suited to describe it. This part of the program has been written in a powerful LISP-embedded problem-solving language called CONNIVER, and is currently running in two segments on a 96K words PDP-10 computer.

The task of the first part of the program is to convert the coordinate information supplied by the touch-sensitive tablet into position information. A position can be loosely defined as "a group of blocks that the subject considers as belonging together at a certain point in time". Certain positions, like the initial position or the balance position, are well-defined, whereas others like the final position or the "intermediate position", which shows up in more sophisticated strategies, are less clear. The program must use strategy information to resolve ambiguities.

The second part of the program has to do with the identification of strategies. The analysis of seriation protocols in terms of strategies has been presented elsewhere. From a programming standpoint, interestingly enough, it turned out to be convenient to re-group the strategies into "stages", since the occurrence of certain actions of the child enables the program to eliminate all the strategies associated with a particular stage.

Each strategy is characterized in three ways:

(a) criteria that must be fulfilled for the strategy to be possible, (b) actions that are sufficient for the strategy to be rejected, and (c) a step-by-step description of the actions that must be performed for the strategy to be accepted. Various "filters" have been introduced to eliminate minor deviations from a basic strategy.

The program performs very well with children of stages one and three, but poorly on the less understood intermediate stage.

"Can a computer replace a developmental psychologist?"
We would be tempted to answer: - "Sometimes...".

TEACHING SERIATION TO PRE-OPERATIONAL CHILDREN:

AN INFORMATION-PROCESSING APPROACH

Jean Gascon and Francoise Gileau

Université de Montréal

The purpose of this experiment is to test the possibility of training children to seriate correctly by teaching them rules assumed to be present in children at the operational level of intellectual development. The rules to be taught are derived from a detailed information-processing analysis of how children at various developmental levels perform on Piaget's weight seriation task (Piaget & Inhelder, 1941). The analysis leads to a formalism called B.G. (Baylor & Gascon, 1974), in which strategies are described in terms of specific production rules of the form (condition \Rightarrow action). If such an analysis is correct, and if a strategy can be taught by training Ss on its separate constituents, - as postulated by Gagne's (1962, 1966, 1970) cumulative learning model - then it should be possible to train pre-operational children by teaching them the B.G. rules for seriation.

The strategy selected for training is one that is often used by the nine and ten-year olds, namely, the "find the heaviest" strategy. It consists of finding, first, the

TABLE I

NUMBER OF Ss GOING FROM STAGE "N" ON
PRE-TEST TO STAGE "M" ON POST-TEST
(N=16)

GROUPS :		EXPERIMENTAL			CONTROL		
From/To: Stage		1	2	3	1	2	3
1	WEIGHT	0	0	4	3	1	0
	LIGHT	0	0	4	4	0	0
2	WEIGHT	0	1	3	0	4	0
	LIGHT	0	1	3	0	4	0

Since all sessions were vidio-tape recorded, it was possible to analyse Ss' performance in a more detailed fashion. These analyses show, without any doubt, that children did learn seriation and could resist counter-suggestions. The analysis also reveals that the rules were not all equally important and that many children discovered some rules by insight.

These results are interpreted as supporting both the incremental learning theory and the information-processing analysis. Implications for Piaget's theory of cognitive development are also discussed.

FORESIGHT AND HINDSIGHT IN AN INFORMATION
PROCESSING MODEL OF THE DEVELOPMENT OF
SERIATION

Gisele Lemoyne and George W. Baylor
Université de Montréal

An information-processing model of aspects of the development of seriation is presented. It is based on a longitudinal study of one child (age 7;2 in April, 1973) on seriation, anticipation, and memory tasks and predicated on Piaget's observation that the child's capacities of foresight (anticipation) and hindsight (rétroaction) are crucial to an understanding of the causal mechanisms of development. The model was empirically derived from two sorts of data: S's protocols and certain complementary experiments on the same three tasks that were introduced to probe further determinants of her behavior that are not revealed by protocol analysis. Some of the data were set aside as a criterion task to permit an assessment of the model's concurrent validity.

S was tested in two different phases: Phase I from April to June, 1973; Phase II from February to March, 1974. During Phase I, S succeeded length seriation (at Stage 3), failed weight seriation (Stage 1), could not anticipate more

than one transformation, nor remember more than two weight or length relations. The complementary probes had little effect of S's performance on the main problems.

In Phase II, approximately nine months later, S attained Stage 3 on weight seriation, was now capable of anticipating two transformations, and of retaining three weight and length relations. During this phase of the experimentation, the complementary problems clearly (and unexpectedly) facilitated S's progress from Stage 2 to Stage 3 on weight seriation.

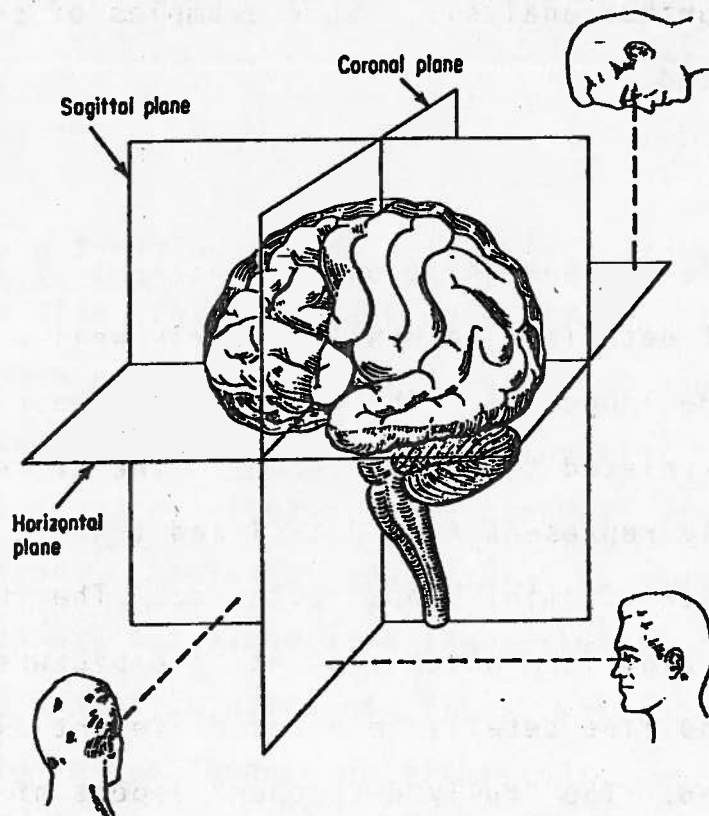
The computer program, written in UCILISP, reproduces S's behavior on the seriation and anticipation protocols. The same strategies generalize across tasks and simulate the task-specific behaviors when the parameters of the program are keyed to the specific task contents. The program shows how a child who does not anticipate more than one transformation, nor retain more than two relations, cannot succeed weight seriation (because of its task demands), nor learn from certain kinds of exercises. With an increase in the program's anticipatory and retentive powers, however, it is able, like the child it simulates, to construct an initial weight seriation strategy (at the intermediate Stage 2 level) and, more importantly, to make use of the complementary problems to emend this strategy to the Stage 3 operational level.

How? In the model a series is defined by the relations that S expects (anticipates) to find between adjacent elements of the series since these anticipations determine his subsequent actions. Thus, the effect of learning is to modify the contents of the anticipations in long-term memory, when possible. In Phase II, before the complementary experiments, an element is thought to be seriated if it is lighter (or heavier) than its neighbor to the left (or right) since a single relation is used to place an element in the series. After the learning engendered by the complementary experiments, it is realized that an element is seriated only if it is at the same time lighter than one of its neighbors and heavier than the other (two relations), which implies the capacity to anticipate two successive comparisons, i.e., two transformations.

The results of the anticipation experiments confirmed that S had the capacity to anticipate two transformations during Phase II. At the same time, it should be added, the complementary anticipation exercises were ineffective in inducing any development in the capacity to anticipate. It remains an open question as to whether progress on the anticipatory capacity can be brought about given the

short-term memory spans (or computing power) associated with various age levels.

Finally, the program is able to predict S's behavior on a specially constructed hidden length seriation task with equal and unequal differences between sticks. This serves as a test of the model's validity.



SEGMENTATION OF SINGLE GRAY LEVEL
PICTURES OF GENERAL 3D SCENES

Dr. T. Kasvand

National Research Council of
Canada

Summary

Gray level pictures of general three-dimensional scenes have been segmented and the segments classified into a few categories. This allows the selection of more specific algorithms for further analysis. Some examples of segmentation are presented.

Introduction

Pictures of black and white objects without excessive texture and small details were handled fairly well with a simplified picture language¹. The objects were described as sets of inter-related "blobs" or atoms. The procedure could not properly represent fine detail and texture due to the large number of "mini-blobs" obtained. The remedy appears to be an algorithm which segments the picture into areas of gross and fine detail, to which different algorithms have to be applied. The "fully developed" aspect of this problem is represented by a gray level picture of a fairly general three-dimensional scene. It is full of both gross

and fine details, to which, as usual, meanings can normally be attached only after the objects in the picture have been recognized. For the black and white pictures, this paradox was partially solved by defining a function $r(x,y)$ from which "blob" centers and limits were computed. In gray level pictures of three-dimensional scenes, this paradox is more difficult to outmanoeuvre, but without proper classification of the gross characteristics of the differing areas in the picture, any "higher level" language will be overloaded by profusion of irrelevant details.

Problem Definition

A single gray level picture of an arbitrary three dimensional scene is assumed given. No a priori information regarding picture content is available. Colours, and temporal or stereoscopic effects available in suitable picture pairs are not considered at present. It is a typical problem we encounter for example when looking at pictures in a newspaper.

Some Steps Towards a Possible Solution

As defined, the problem is quite formidable. Without some type of a "world model", this problem, of course, cannot be solved. The model will, of necessity, have to be made up of representations of object segments and their interrelationships in space and time. The picture segmentation process, however, will have to be started under the conditions given in the problem definition.

Step 1:

The picture $z = f(x,y)$ is treated as one global unit or segments, where x and y represent picture coordinates and z the gray levels. A small (3×3 or 5×5) local area is used to compute function like average, ∇ , ∇^2 , etc., and combinations of these resulting in a new function $V(x,y)$ defined over the picture domain. In principle $V(x,y)$ represents also matched filtering, mask matching, etc., but since picture content is still entirely unknown, $V(x,y)$ should be rather general and preferably independent of both size and orientation.

Step 2:

For a rather general $V(x,y)$, the pictures have been segmented according to two methods. The first method (I) is based on the amplitude histogram of $V(x,y)$, the other (II) on the local extremum values of $V(x,y)$.

Method (I) consists of the following:

- (a) Form amplitude histogram of $V(x,y)$.
- (b) Divide the histogram into high, middle and low amplitude portions.
- (c) Generate a binary mark corresponding to each of the histogram portions and smooth it.
- (d) Display the picture through this mask to see that portions of the picture have been selected.

Method (II) consists of:

- (a) Find the local extrema of $V(x,y)$.
- (b) Form little neighbourhoods around the extrema, which serve as binary masks.
- (c) Same as (d) in Method (I).

Steps 3,4:

The following steps will be increasingly dependent upon the form of $V(x,y)$ and the type of picture segment chosen.^{2,3}

Some Types of $V(x,y)$ Which Have Been Studied Experimentally

1. $V(x,y) = f(x,y)$, i.e. the gray level values in picture.
The extreme ends of the gray level histogram give the very bright and the very dark regions in the picture. The very bright regions may be the light sources, reflections from smooth surfaces or exceptionally white areas. The dark regions may represent "cavities" into which light does not penetrate, shadows or very dark objects.



2. $V(x,y)$ = A Form of Spatial Differential

The following algorithms have been tried, where
 n = the number of picture elements in the local area
 (L) , ($L = 3 \times 3$ or 5×5), $/./$ = magnitude of, ∇ = spatial
gradient.

$$(a) \quad V(x,y) = \frac{1}{n} \left(\sum_n /f(x,y) - \frac{1}{n} \sum_n f(x,y) \right)$$

$$(b) \quad V(x,y) = \text{standard deviation of } f(x,y) \text{ in } L$$

$$(c) \quad V(x,y) = / \nabla f(x,y) / \text{ in } L$$

$$(d) \quad V(x,y) = / \nabla^2 f(x,y) / \text{ in } L$$

$$(e) \quad V(x,y) = / \text{angle difference of } \nabla f(x,y) / \text{ in } L$$

$$(f) \quad V(x,y) = \text{density of contour elements in } L$$

$$(g) \quad V(x,y) = / \text{slope difference of contour elements} / \text{ in } L$$

All these $V(x,y)$ are independent of orientation but not
of size of objects. For gray level pictures algorithms
(a) to (d) produce fairly consistent results, (e) has
a bias against straight line data, (f) and (g) are being
tried. (On certain types of line drawings (f) produced
fairly interesting results.) In general, the high end
of the histogram of $V(x,y)$ brings out areas of densely
packed highly contrasting detail (i.e. the "busy" areas).
The low end gives rather featureless or "quiet" areas.
The middle region represents the "residue" to which the
algorithm may be applied again for further segmentation.

Conclusion:

The $V(x,y)$ function gives a crude classification of the various areas in the picture. A variety of such functions should refine the classification to a certain degree. This preliminary classification of the segments can be used to call upon algorithms which are more suited for further analysis of given areas in the picture. It would be very desirable to have a standard set of pictures available which typify various scenes.

Appendix: An Example of Segmentation

General: The original picture is quantized to a 256×256 matrix of up to 256 gray levels, originating from Professor W.K. Pratt, Image Processing Institute, University of Southern California. Each processing step adds a short commentary to the picture data (probably illegible after reproduction). The density modulated histogram (closest to the picture) shows the gray level distribution, where local histogram brightness corresponds to brightness in the display. The other is the variability histogram, with selected busy, mid and quiet ranges shown. Fig. 1 shows the original picture, Fig. 2a (manually) extracted area (head) and Fig. 3 the variability $V(x,y) = \text{angle difference of } \nabla f(x,y) / (\text{form } e)$, displayed in "3D" with top right corner (of Fig. 1) lowest.

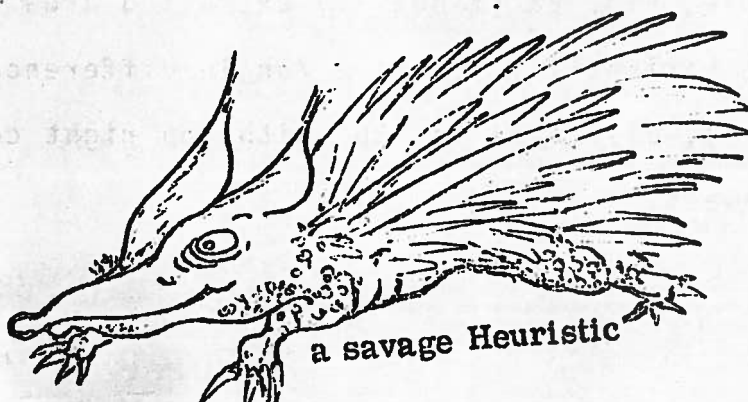
Method I Applied to the Whole Picture: 20% of the "busy" and 20% of the "quiet" ends of the variability histogram were selected. Fig. 4 shows the "busy" parts of the picture, Fig. 5 the intermediate range of Fig. 6 the "quiet" parts. Figures 7 and 8 show the "busy" segments for 10% and 30%, respectively.

Method I Applied to the Head Alone: Selecting again the 20% limits gives Fig. 9 as the "busy", Fig. 10 as the intermediate and Fig. 11 as the "quiet" parts of the head.

Method II Applied to the Whole Picture: 64 of the highest local maxima in $V(x,y)$ were selected and expanded to 5×5 matrices (Fig. 12). Fig. 13 shows these maxima in context of the original picture.

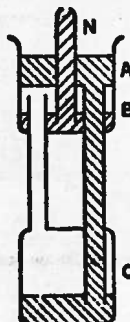
Comments:

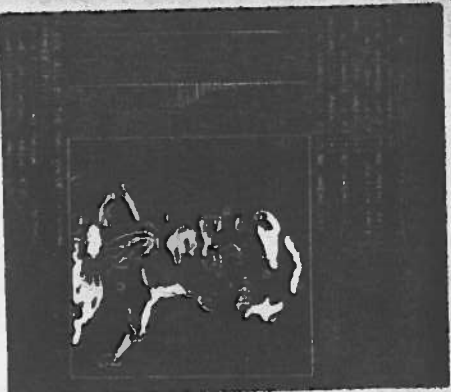
The procedure always forces a segmentation. The results derived from the tapering and (usually the "busy" end) of the histogram are fairly consistent. Sometimes the segments appear directly usable as elements in a "picture language".



REFERENCES

1. **Graphic Languages**, pp. 163-186, Editor: F. Nake & A. Rosenfeld, North-Holland Publishing Co., 1972.
2. **Scene Analysis Using a Semantic Base for Region Growing**, Y. Yakimovsky, Stanford Artificial Intelligence Lab., Memo AIM-209, June 1973.
3. **Textural Features for Image Classification**, Robert M. Haralick, K. Shanmugam, and Its'hak Dinstein, IEEE Transactions, Vol. SMC-3, No. 6, pp. 610-621, November 1973.

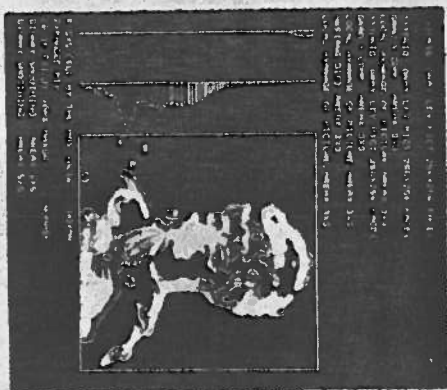




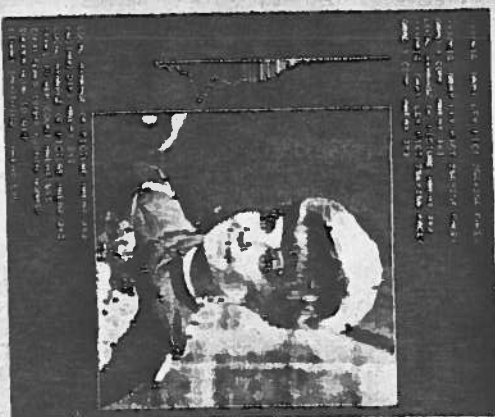
4



7



8



13



5



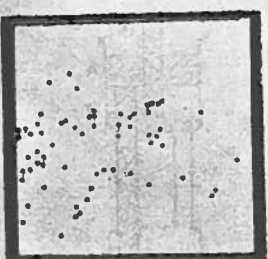
1



3



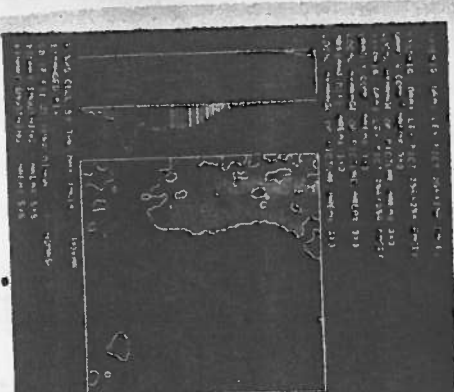
2



12

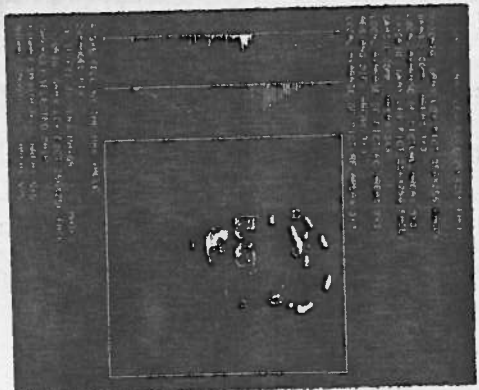
10

11



6

9



FIGURES
1 TO 13

**AN EXTENSIBLE FEATURE-BASED PROCEDURAL QUESTION ANSWERING
SYSTEM TO HANDLE THE WRITTEN SECTION OF THE BRITISH COLUMBIA
DRIVER'S EXAMINATION**

G. McCalla and M. Kuttner

Department of Computer Science

University of British Columbia

Vancouver, British Columbia

Abstract:

An extensible feature-based procedural question answering system to handle the written section of the British Columbia driver's examination is proposed. Its capabilities are discussed and suggestions are made for future research.

Keywords: kuntji, palabra clave, slagwoord, ufunfuo,
 schlissel, clavis, nøgle, stichwort, gi'ghbank

The Paper:

For nearly as long as there have been motor-cars, the language associated with them has been subjected to intense scrutiny. Early papers, such as McSpud's (1908) landmark treatise, tended to take a rather philosophical outlook, and it wasn't until Marmaladoff's (1961) analysis that any rigorous linguistic description was attempted. Although not of great import in itself, this paper did spur the use of more formal methods, and eventually Ste. Couche (1967) and his group brought

into play the full power of the electronic computer. Two of Ste. Couche's students, Madatter and Marchaire (1969), went still further, introducing A.I. methodology into the discipline. The model they developed was somewhat limited in its domain of discourse (handling only one question), but it did serve as the starting point for many other investigations, including our own. The reader who would like a more complete description of this active research area, is referred to the excellent survey done by ffoulkes-Barr (1972). -

The computer model presented here (we call it D-FENDER) is an expanded version of the HUB-CAP system described in our earlier paper (McCalla and Kuttner (1984)). Although retaining the basic flavour of HUB-CAP, the new system is more robust, not running amuck when presented with sentences longer than 3 words, and giving the correct answer almost 80 percent of the time (an increase of close to 80 percent). This startling improvement was achieved using what promises to be a new fundamental axiom of A.I.: the ignorance principle, which states, essentially, "why clutter up the system with a lot of knowledge?" By being totally devoid of knowledge, any system is able to overcome combinatorial explosion and the related dilemma of uncontrolled back-up, there being nothing to explode and no where to back-up to. No doubt, many other thorny A.I. problems will turn out to

be equally amenable to an ignorance solution, but exactly which problems, and exactly how the principle would apply, are questions that remain to be answered.

D-FENDER is the first true ignorance model(IM), embodying as it does the ignorance principle in every stage of its processing. Semantics, pragmatics, and syntax have all been completely eliminated (thus resolving many contentious linguistic issues). Instead, each input utterance is tested for the presence of certain features or keywords which are themselves programs. If a keyword is found, its associated procedure is executed to (perhaps) perform further tests and to produce the appropriate response.

Implementing these ideas in a computer model has been no trivial task, and literally several man-minutes have gone into constructing, debugging, and testing the D-FENDER program. A structured program, it is written in LISP, uses 16K bytes of core on an IBM 370/168, and runs interactively under the MTS operating system. A sample run is shown in Figure 1.

```
* (HIGHWAY-TEST)
> TYPE IN THE QUESTION SURROUNDED BY BRACKETS.
> WHEN YOU WANT TO STOP, TYPE IN (GOODBYE).
> ENTER YOUR QUESTION NOW, PLEASE.
* 1. (MUST EVERY TRAILER WHICH, OWING TO SIZE OR
* CONSTRUCTION TENDS TO PREVENT A DRIVING
* SIGNAL GIVEN BY THE DRIVER OF THE TOWING
* VEHICLE FROM BEING SEEN BY THE DRIVER OF THE
* OVERTAKING VEHICLE BE EQUIPPED WITH AN
* APPROVED MECHANICAL OR ELECTRICAL SIGNALLING
* DEVICE CONTROLLED BY THE DRIVER OF THE
* TOWING VEHICLE)
> THE ANSWER IS:
> YES
> ENTER YOUR QUESTION NOW, PLEASE.
* 2. (HOW MANY FEET PER SECOND ARE YOU TRAVELLING AT
* 60 MILES PER HOUR)
> THE ANSWER IS:
> 88 FEET
> ENTER YOUR QUESTION NOW, PLEASE.
* 3. (TO WHAT MUST THE COUPLING DEVICE BETWEEN A
* MOTOR-VEHICLE AND TRAILER BE AFFIXED)
> THE ANSWER IS:
> YES
> ENTER YOUR QUESTION NOW, PLEASE.
* 4. (WHERE MUST THE REAR CLEARANCE LAMPS BE MOUNTED)
> THE ANSWER IS:
> EXTREME RIGHT AND LEFT SIDES
> ENTER YOUR QUESTION NOW, PLEASE.
* 5. (WHERE THE SERIAL NUMBER OF A MOTOR-VEHICLE HAS
* BECOME ILLEGIBLE OR HAS BEEN REMOVED OR
* OBLITERATED, WHAT IS NECESSARY)
> THE ANSWER IS:
> IMMEDIATELY NOTIFY THE SUPERINTENDENT OF MOTOR-VEHICLES
> ENTER YOUR QUESTION NOW, PLEASE.
* 6. (THANKS AND GOODBYE)
> AU-REVOIR ET BONNE CHANCE
```

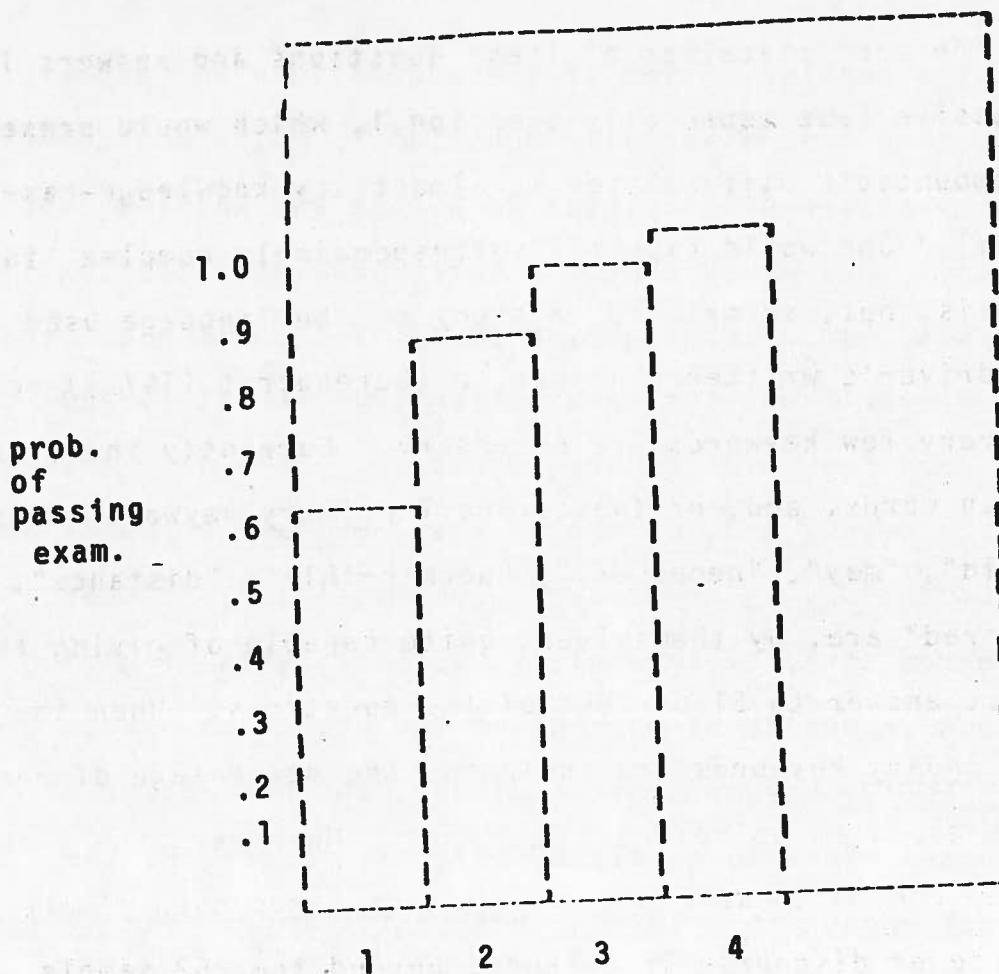
Note: Input to the program is preceded by * .

Output from the program is preceded by > .

Question numbers have been added later.

Figure 1 - Sample Run

The sophistication of these questions and answers is impressive (see especially question 1, which would present insurmountable difficulties to almost any knowledge-based system). One would expect a correspondingly complex linguistic analysis, but, surprisingly, a study of the language used in the B.C. driver's written examination (Bureaucrat (1972)) reveals that very few keywords are necessary. Currently there are only 30 such words, and, of these, the 7 primary keywords "must", "should", "may", "necessary", "permissible", "distance", and "required" are, by themselves, quite capable of giving the correct answer to 51 percent of the questions. When the other 23 secondary keywords are included, the percentage of correct responses rises to nearly 80 percent. Moreover, the program is general enough to maintain these percentages even if the universe of discourse is extended beyond the 262 sample questions in Bureaucrat. A statistical evaluation (see Graph 1) shows that, with an alpha error of .05, at most 4 attempts will be needed for the program to pass the examination (i.e. be correct on at least 16 of 20 selected questions). When it is realized that the average human needs 3 or 4 trials before achieving success (a ball-park estimate), it becomes abundantly clear how intelligent the model really is.



no. of attempts

GRAPH 1 - Statistics

Nevertheless, the model does have its flaws. For example, the answer to question 3 of the sample dialogue should be "FRAME", not "YES" (the frame problem); and the system is weak in the motor-cycle sub-domain. Difficult as they are, though, these problems are all almost certainly amenable to ad hoc solutions, in keeping with the ignorance principle. D-FENDER might also be indicted for being rather narrow in its choice of subject area, the kinds of sentences it processes, and so on. However, there is no a priori reason to believe that the system could not be expanded to handle other kinds of tests; that is, there could be an IM to write pilot's licence written exams, an IM to pass Ph.D. qualifying exams, or an IM to participate in a thesis defence. In fact, THESIS D-FENDER is now being constructed.

Eventually, we would like to apply the ignorance principle to other areas of A.I. altogether. Thus, an ignorance theorem prover (ITP), an ignorance game player (IGP), or even an ignorance robot (IR) could be devised, undertakings which would be immeasurably aided by the vast amount of relevant research already in existence. Whatever the fate of these plans, however, the success of D-FENDER in its own domain is undeniable. It is stunning proof of the value of not knowing what is going on.

REFERENCES

- Bureaucrat (1972). Anonymous Bureaucrat, British Columbia Guide to Safe Driving, Dept. of Highways, Government of B.C., Victoria, B.C., 1972.
- ffoulkes-Barr (1972). Clive Perkin ffoulkes-Barr, "Natural Language and the Driver", Playcar, September, 1972, p. 42.
- Madatter and Marchaire (1969). Thrumboldt Madatter and Guse Steppe Marchaire, "Heuristically Proving Theorems About Mechanistic Discourse", Com. S.P.C.A., 39, March 1969, pp. 23-45.
- Marmaladoff (1961). Djamiatin Marmaladoff, "NP + VP = MC²: An Arithmetic Approach to Relativization in Honda", Int. J. Comp. Sec. Ling., 4, July 1961, pp. 321-4321.
- McCalla and Kuttner (1984). Gordon McCalla and Michael Kuttner, HUB-CAP: How to Undertake Back-up in Cars by the Application of Procedures, Tech. Doc. 1, Toy Department, Centre for Over-Exposure and Sun Worship, Wreck Beach, B.C., November, 1984.
- McSpud (1908). Gruncheon McSpud, Punctuation Marks in the Motor-Vehicle Statutes of the Prairie Provinces, Wine Press, Grasshopper, Saskatchewan, 1908.
- Ste. Couche (1967). Jean-Loup Ste. Couche, "Driver Training from a Data Processing Point of View" in L'Ordinateur, La Voiture et Le Chauffeur, Roger T. Saurus, ed., Pamplémousse, Poisson, and Rutabaga, Ltd., Oookpik, Quebec, 1967.

**A LANGUAGE IN WHICH TO EXPRESS QUALITATIVE
SCIENTIFIC KNOWLEDGE**

D.R. Skuce

**Department of Electrical Engineering
and
Montreal Neurological
Institute
McGill University**

KAL (Knowledge Acquisition Language) is a language which is intended to be used by scientists working in domains whose facts are usually not amenable to explication via traditional mathematical formalisms. The most obvious example of such a domain is biology, in which only negligibly few micro-domains have in any way benefitted from what is usually called "mathematical modelling". The vast proportion of biological knowledge is expressed in technical natural language, and the models remain purely a fuzzy mental agglomeration of words, rules and images, as different as the minds of scientists.

A language such as KAL, with precise syntax, and particularly, semantics, would be of great benefit even as a tool for interpersonal communication, and as a medium for verifying assumptions and deductions. In addition, however, an entirely new prospect arises should it be possible to implement a machine which "understands" the language in the normal AI sense. One might then enter hopefully all the relevant facts of some domain, and interact with the machine in performing inductive and deductive inferences, predicting experimental outcomes, and discovering underlying principles.

KAL sentences are presently of two kinds: the "proposition", and the "implication". The proposition is a declarative active or stative simple English-like sentence: NP VP, where NP denotes the subject set, and VP has the familiar case structure, using keywords for the cases which employ additional NP's. For example:

THE HANDBRAKE IS-CONNECTED TO THE REAR WHEELS
BY-MEANS-OF A CABLE LOC-UNDER THE FLOOR

The underlined words are keywords. Additionally, a proposition may specify a constraint on the possible values of the subject-verb pair viewed as a functional application with the verb denoting the function, and the subject the argument, as in:

EACH PERSON'S SEX-IS ONE-OF (MALE FEMALE)
EACH ADULT'S AGE-IS GT (20)

If subsequent facts are entered which enlarge, restrict or contradict such constraints, the user is to be warned.

The implication is an IF-THEN pair of propositions, as in:

IF THE CABLE OF THE BRAKE OF A CAR IS JAMMED
THEN THAT BRAKE IS INOPERATIVE.

It is important to note that the non-underlined words are vocabulary introduced at will by the user. The user must keep account of how each word, particularly verbs, is being used (its distribution), as this is the only meaning the machine has for it. The syntax is sufficiently rigid to eliminate nearly all ambiguities. If one occurs, the user is interrogated.

The NP may contain 1st and 2nd order quantifiers, as in:

EACH PERSON HAS THEIR-OWN SET-OF FRIENDS

and may include adjectives, prepositional phrases and elementary subordinate clauses, as in:

THE BIG RED BALL ON THE TABLE BELONGS-TO
THE MAN WHO BOUGHT THE HOUSE OF JOHN

An additional important concept familiar to programmers as a data structure, is the notion of whole-part relations, or 1-1 correspondence. These are usually either physical, as in:

EACH HAND IS-PHYS-PART-OF ITS ARM

or conceptual, as in:

EACH MARRIED-COUPLE CONCEPTUALLY CONSISTS-OF
(A MAN, A WOMAN)

The 1-1 correspondences of these structures are automatically maintained, so that if we say:

BILL HAS-AS-PHYS-PART AN ARM

then the question:

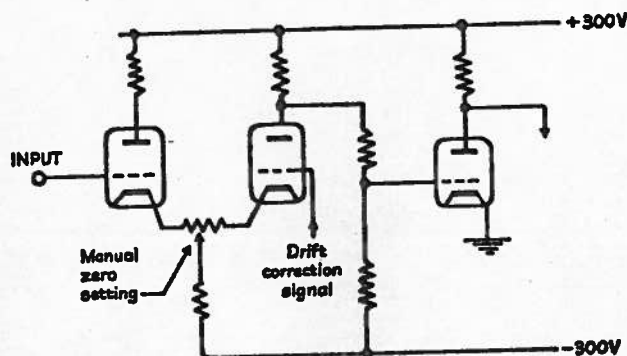
BILL HAS-AS-PHYS-PART A HAND?

is answerable without explicit implication statements, by by tracing HAS-AS-PHYS-PART pointers, just as set membership has special pointers to deduce inherited properties.

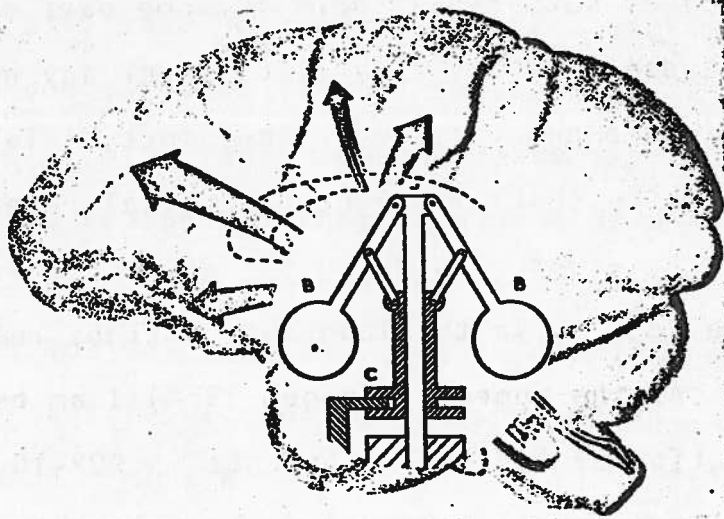
Time, change and actions are notoriously difficult problems for such systems. An attempt is now being made to introduce the notions of state, event (a sudden change of state) and process (an on-going, usually smoother change of state), which clearly are necessary in the description of most real world situations. Since the implementation is being attempted in CONNIVER, some of the usual "ADD-REMOVE" methodology may be used for events. The problem of processes is much more difficult. One requires some discrete representation of the passage of time, and the division of concepts into functionally related "black boxes", as is done in conventional simulation languages such as SIMSCRIPT, GPSS SIMULA, etc. One must then introduce an arithmetic capability. The brute force approach would be to treat each clock interval

as an event, and run ADD-REMOVE routines to update all changing variables. The inefficiency would be horrendous.

Even without a complete computer implementation, a precise definition of KAL on paper would be sufficient to test the adequacy of the available constructs, by hand checking the answering process for some set of questions chosen to be as "meaningful" as possible. An important aspect of KAL is that it be actually usable (i.e. comprehensible) by such researchers in some part of their work, though I don't expect they'll discover any new results at this elementary stage. Therefore, I intend to try to teach KAL to several bright biological researchers, and with them, to attempt to describe some micro-world of their technical domain in the language, thereby refining the language. At the moment (October 1974) I am beginning the implementation of the basic concepts in PDP-10 CONNIVER, and am preparing a manual intended to be comprehensible to biological scientists and physicians.



RESEARCH ABSTRACTS



A PROTOTYPE SPEECH RECOGNITION SYSTEM

by James F. Allen

University of Toronto

A system for processing and recognizing sentences of carefully spoken continuous speech is described. The entire system is driven from the syntactic level by a parser for an Augmented Transition Network grammar. Other features of interest in this system include the use of the syllable as the basic unit of recognition, the forming of word patterns interactively and the fast retrieval of words from the dictionary by sound characteristics. The parser has the facility to use Prosodic Information as a guide in deciding which arcs are the most likely to succeed, and uses co-routines with shared data structures to allow multiple parse paths to be pursued at any one time.

This thesis is soon to be published as a University of Toronto Technical Report.

USE OF BOOLEAN TREE FUNCTIONS TO PERFORM HIGH-SPEED PATTERN CLASSIFICATION AND RELATED TASKS

W. W. Armstrong and G. Godbout

Departement d'informatique

Universite de Montreal

In this working paper, a new approach to high-speed pattern classification is described which uses combinational circuits formed of two-input, one-output elements interconnected in the form of a binary tree. The functions so realized are shown to be, on average, very insensitive to perturbations of input variables, and so can perform "smooth" extrapolation from training data to test data. A theory is developed and computer simulation experiments are described involving recognition of handwritten numerals, numerical taxonomy, and medical diagnosis.

PROPERTIES OF BOOLEAN FUNCTIONS WITH A TREE DECOMPOSITION

G. V. Bochmann and W. W. Armstrong

Departement d'informatique

Universite de Montreal

Boolean functions that have a multiple disjoint decomposition scheme in the form of a tree are considered. Properties of such functions are given for the case that the functions are increasing, unate, and/or have no vacuous variables. The functions with a binary decomposition scheme are of special interest. The modulus of sensitivity is defined, and evaluated for some classes of functions. It is found that the sensitivity for the class of functions with a given disjoint binary decomposition scheme is much smaller than for the unrestricted class of boolean functions.

(A shortened version has appeared in BIT 14, pp. 1-13, 1974).

A CONVERGENCE THEOREM FOR LOGICAL NETWORK ADAPTATION

W. W. Armstrong and G. V. Bochmann

September 1972

Departement d'informatique

Universite de Montreal

A task of classifying n -dimensional boolean vectors into two classes may be described by a boolean function of n variables. When n is large, finding an economical realization (even approximate) is difficult. In this article a method of synthesis is described whereby one first fixes a composition scheme in the form of a binary tree, and then determines by an adaptive algorithm the functions of two variables to be inserted at the nodes of the tree. The basic problem is to make good use of the given tree. A theorem is proved giving conditions under which the algorithm leads to an exact realization. The question of approximate realizations under the most general conditions is raised in order to appraise the significance of the theorem.

EXPERIMENTS IN SERIATION WITH CHILDREN: TOWARDS AN INFORMATION
PROCESSING EXPLANATION OF THE HORIZONTAL DECALAGE

George W. Baylor and Gisele Lemoyne

Universite de Montreal

Thirty-seven children between the ages of five and ten were tested on three seriation tasks: length, weight, and a specially constructed "hidden length" seriation designed to eliminate the typically observed horizontal decalage between the length and weight tasks. The experimental results confirmed the main hypotheses: that the hidden length seriation was more difficult than the normal length seriation and of approximately equal difficulty to the weight seriation. An information processing analysis was then carried out with one child at the operational level of cognitive development, Pierre. This led to the construction of a performance model, a computer program cast as a set of production rules, that simulated in detail Pierre's behavior on the three tasks. The general and task specific elements of the program show how the length seriations are facilitated by perceptual factors and the organization of memory, though within the framework of a common insertion strategy that Pierre employed on all three tasks.

(To appear in Canadian Journal of Behavioral Science, 1975).

SEMANTICS OF ENGLISH IMPERATIVE SENTENCES FOR SIMPLE
ARITHMETIC USING A RECURSIVE AUGMENTED TRANSITION

NETWORK GRAMMAR

William J. Blewett

University of Western Ontario

The purpose of this thesis is to investigate some of the problems of natural language "understanding" in computer systems. It was decided to implement such a system for a particular conversational domain, and as the problem of understanding language falls naturally into several subproblems, to use the best available methodologies for each sub-system. As it turned out, the choice of conversational domain influenced the design somewhat, the major effect being the elimination of the need

for an explicit knowledge base of semantic information. However, this aspect allowed us to ignore the problems involved in the structure and manipulation of knowledge bases and to concentrate more on the analytic processes involved in extracting the "meaning" of linguistic utterances. The methodologies called upon were:

1. augmented transition networks for parsing
2. Winograds methodology for the evaluation of noun phrases
3. the procedural representation of meaning
4. case grammar as a deep structure representation of sentences.

The implementation was written in POP-2. The parser was translated (and slightly modified) from a LISP program obtained from Dr. R. Reiter at the University of British Columbia. The grammar also borrows substantially from the work of Jean Jervis of U. B.C.

The major factor involved in choosing simple arithmetic as the domain of discourse for this natural language system was a decision to avoid, if possible, the deep problems involved in maintaining a data-base. We wished to concentrate on the analysis of language: that is, on the identification and disambiguation (both syntactic and semantic) of linguistic fragments, rather than on the interaction of deep-structures with a data-base to produce a response (or answer) to statements. Once the deep structure of a statement has been deduced, we wish to limit the amount of computation necessary to produce a response.

The syntactic and semantic structure of English sentences about arithmetic can remain as complex as we wish them to be. However, having once deduced the deep structure of these statements, it is usually possible to map them into simple computer procedures for the calculation of answers.

The intention of the research was to investigate the process of disambiguating linguistic structures such as sentences and phrases. The conversational domain, though limited, still contains a large variety of forms of ambiguity. Some words may belong to more than one syntactic category such as "prime" which can be both a noun and an adjective.

A word in one syntactic category may have more than one definition in that category. The syntax of a sentence or phrase may be ambiguous with respect to the grammar and even a grammatically correct structure may be meaningless with respect to the semantics of the conversational domain.

Our main contribution to this system lies in the semantic procedures of arithmetic which evaluate words, phrases and clauses in order to disambiguate their structure and meaning. Also we have suggested a method of handling conjunction in English by a modification of the recursive transition network parser. (M.Sc. Thesis, 1974).

PROBLEMS OF UNDERSTANDING SENTENCES BY COMPUTER

Alexander T. Borgida

University of Toronto

The present research involves the semantic interpretation of English sentences, with special application to question-answering systems.

The attack on the problem is multi-dimensional although we have segmented it, for convenience, into modules such as morphology, syntax, predicate semantics, etc.

We first obtain a form of syntactic deep structure for the input sentence, using an Augmented Transition Network parser which we developed. The purpose of the resulting "parse tree" is to guide us to those pairs of "words" in the sentence which should be directly linked semantically.

These semantic links are then discovered using sophisticated case frames for predicates, some heuristics and a semantic network containing general world knowledge as well as the immediate context in which the sentence was uttered.

Two particular points of interest to us are multiple adverbial (prepositional phrase and adverb) modification and the complex noun phrase.

For an adverbial, we get involved in an active search for the constituent which it is most likely to modify, using case frames and heuristics based on the special parsing strategies developed for adverbials in the A.T.N.

In the case of the complex noun phrase, with each type of modifier (i.e. adjective, possessive, participles, relative clause and prepositional phrases) we associate a list of alternatives to be checked in order to establish the meaning relationship between the modifier and modified head noun. Each step in such a list usually involves a look at the semantic network and, when it succeeds, results in either the identification of the referents of the modifier and modified in terms of the context, or, the creation of a link between them, possibly through some "understood" event, which explicates their relation. (e.g. "a good knife" is understood to mean "a knife which one can cut with well").

Future research will delve into noun-noun modification and will probably involve the creation of a sufficiently large and well-structured semantic network on which one can test the heuristics developed. (Master's Thesis, University of Toronto, 1974).

A PLANE EUCLIDEAN GEOMETRY MACHINE

Patrick Deer

University of Western Ontario

The recent development of goal-oriented languages, such as PLANNER and CONNIVER, has provided a high-level, "natural" formalism for representing mathematical knowledge in the form of programs. These languages facilitate backtracking, versatile pattern matching, invocation of appropriate procedures by matching the target pattern, etc.

A geometry machine for a limited area of plane Euclidean geometry is built by employing these pioneer programming facilities available in PLANNER. This machine is able to solve many of the geometry problems presented to it in an appropriate mathematical description. This project is an implementation of the theoretical concepts outlined by Mr. Ira Goldstein (1973) in his paper "Elementary Geometry Theorem Proving". (M.Sc. Project Report, 1974).

ON THE INTERPRETATION OF DRAWINGS AS THREE-DIMENSIONAL SCENES

Alan Keith Mackworth

University of British Columbia

This thesis is concerned with scene analysis which is in the vision area of artificial intelligence research. Scene analysis is characterized by computer programs that interpret pictures as three-dimensional scenes. It is shown that considerable progress has been achieved by restricting the pictures to line drawings, possibly derived from grey-scale representations, of polyhedral scenes. The usual strategy for such a program is to abstract cues from the picture which address a repertoire of models which, in turn, address the picture. The cues are mostly uninterpreted picture fragments while the model repertoire consists of a list of polyhedral prototypes. It is suggested that the cues be enriched by interpreting the picture fragments and the models be enhanced by allowing models to be less monolithic than complete object prototypes. To achieve this end, particular attention is paid to the organization of the visual world into surfaces and their relationships. By exploring the properties of a representation of surface and edge orientations it is demonstrated that there are general coherence rules that the surfaces and edges must satisfy. A program POLY, that achieves the interpretation of line drawings is described. POLY exploits those coherence rules thereby dispensing with the predetermined interpretations of categories of picture junctions as corners used by previous prototype-free edge-labelling procedures. An annotated listing of POLY and some typical results are provided. Several possible extensions to this method are considered with emphasis placed on the integration of models of surfaces and their relationships into the program. There follows an examination of the conceptual and practical connections between POLY and the programs discussed earlier in the thesis. A brief contrast between theories of human and machine vision reveals complementary weaknesses and strengths. The thesis concludes with a reference to a proposed scene analysis task which will be the context for further development and implementation. (Thesis submitted to the University of Sussex for the degree of Doctor of Philosophy, January 1974).

USING MODELS TO SEE

Alan K. Mackworth

University of British Columbia

Scene analysis programs offer the hope of providing a more adequate account of human competence in interpreting line drawings as polyhedra than do the current psychological theories. This thesis has several aspects. The aspect concentrated on here is that those programs have explored a variety of methods of incorporating a priori knowledge of objects through the use of models. After outlining the range of models used and sketching some psychological theories, the various proposals are contrasted. This discussion leads to two new proposals for exploiting model information that involve elaborations of an existing program, POLY.

MACRO VS MICRO PATTERN DEVELOPMENT

Lawrence J. Mazlack

University of Guelph

The process of constructing or recognizing patterns may proceed on either a macro or micro basis. This paper explores pattern development on both a macro and micro basis as applied to the problem of crossword puzzle construction. Human puzzle constructors apparently proceed on a macro basis. However, it appears that the most successful machine approach is a micro approach. This paper explores the strategies, data structures, and decision structures to develop both macro and micro crossword puzzle patterns.

(Presented at the Second International Joint Conference on Pattern Recognition, Copenhagen, August, 1974.)

THE 2.PAK LANGUAGE: PRIMITIVES FOR AI APPLICATIONS

Lucio F. Mellì

University of Toronto

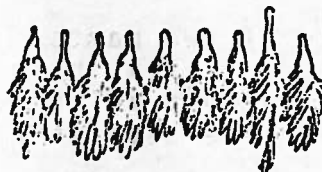
The thesis introduces a language, 2.PAK, whose main aim is to provide a set of primitives suitable for Artificial Intelligence applications. In addition, 2.PAK tries to incorporate principles obtained from research into programming languages in general to aid in program development and the achievement of readable, reliable programs.

The main language features are:

- (a) A wide selection of data types including strings, lists, hash tables, a directed, labelled graph data base, and user defined records.
- (b) A more general control structure in the form of co-routines which can be suspended and later resumed.
- (c) A type of backtracking compatible with the hierarchical control structure and with the coroutine control structure.
- (d) A generalized pattern matcher whose subject for the match is not restricted to a set of language designer decided data types.

Finally, it should be noted that although 2.PAK is not as high level as some of the other AI languages, the thesis shows how one can easily extend the language, by using its abstraction facilities, to obtain such features as: more flexible generators than CONNIVER, a pattern directed function calling mechanism, and list pattern matching primitives.

This thesis is soon to appear as a Technical Report of the Department of Computer Science, University of Toronto. 2.PAK is currently being implemented.



Heuristics attempting
to deal with equality

INFERENCES OVER DATA RECORDS

Peter Wanner

University of Western Ontario

The creation of a large number of related data records always contains the possibility of inconsistencies in the data. It is, therefore, necessary to perform certain tests which find as many of these errors as possible. When the records are created in an interactive mode, one would like to detect the inconsistencies immediately in order to make the corrections before storing the records on disk. This means that in addition to the information about the structure of the records, the user has also to supply a set of assertions which define what "consistency" of the data for a particular class of records is.

Assertional systems like ABSY or ABSET are capable of using assertions not only to verify the consistency of the existing data, but also to infer values to still unevaluated datums. In terms of a record construction system, this would mean that certain record components obtain their values from already existing components by inference from the given assertions.

The purpose of this project was to apply the ideas used in the assertional systems mentioned above and to implement in POP-2 a part of an inferential and interactive record construction system.

(M.Sc. Project, 1973).

ENGLISH GENERATION FROM A SEMANTIC NETWORK IN A QUESTION-ANSWERING ENVIRONMENT

Harry K.T. Wong

University of Toronto

This thesis presents a system which generates English sentences or noun phrases from a semantic network.

The system consists of two parts. The first part, SELECTOR, is responsible for creating a graph by copying a subpart of the semantic network which contains the answers to the user's questions. The next step is to map this

graph to a more surface representation. The second part, the OUTPUT routine, then produces English sentences from the surface structure which the SELECTOR provides, using only syntactic information in its dictionary.

This thesis will soon be available as a Department of Computer Science Technical Report.

THE TORUS PROJECT

The University of Toronto

TORUS (TORonto Understanding System) is a natural language understanding system, being developed at the University of Toronto, which interfaces with ZETA, a relational data base management system. The joint TORUS-ZETA project is designed to be an intelligent query system which will enable a user to communicate via natural language. TORUS makes no presumptions regarding the user's knowledge of the structure and organization of information in the system. It does assume, however, that the user is communicating with a purpose, i.e. to retrieve information about the domain of discourse (which currently is student files.) This "dialogue with a purpose" will, hopefully, narrow the vocabulary and knowledge that the system need possess in order to understand a user's request.

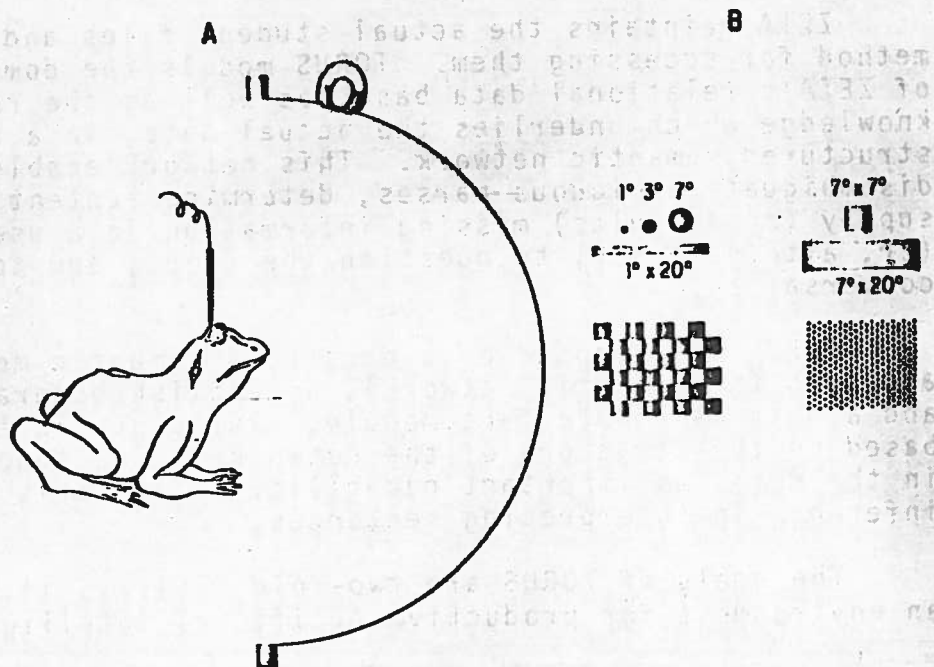
ZETA maintains the actual student files and provides a method for accessing them. TORUS models the domain structure of ZETA's relational data base, as well as the real-world knowledge which underlies the actual data, in a case-grammar structured semantic network. This network enables TORUS to disambiguate ambiguous parses, determine sentential referents, supply (by defaults) missing information in a user's query (or, alternatively, to question the user), and to carry on a conversation.

TORUS is composed of a parser, a semantic module (which accesses the semantic network), an English generation module, and a data base interface module. Inference in the system is based on the structure of the network and on theorems embedded in the net. An important capability of TORUS is its use of inference in interpreting sentences.

The goals of TORUS are two-fold. First, it should provide an environment for productive Artificial Intelligence research.

Secondly, it should be a prototype of a practical system, and thus be convincing evidence that AI can be practically oriented. We hope to have an operational prototype in November which will fulfill these goals. TORUS is being implemented on an IBM s370/165 II in SPITBOL and 1.PAK (an AI language written at the University of Toronto), while ZETA is being implemented in PL/I.

The TORUS project is being developed by the following people: Professor John Mylopoulos, and (in alphabetical order) Alex Borgida, Phil Cohen, Nick Roussopoulos, John Tsotsos, and Harry Wong. The ZETA project is under the direction of Professors Tsichritzis, and Schuster.



LIST OF CONTRIBUTORS

Allen, James F.	Dept. of Computer Science, University of Toronto, Toronto, Ontario, M5S 1A1.
Armstrong, William W.	Dépt. d'Informatique, Université de Montréal, 2101 blvd. Edouard-Montpetit, Montréal, Québec.
Baylor, George	Dept. of Psychology, University of Montreal.
Blewett, William	Dept. of Computer Science, University of Western Ontario, London, Ontario, N6A 3K7
Bochman, G.V.	Dépt d'Informatique, Université de Montréal.
Borgida, Alexander T.	Dept. of Computer Science, University of Toronto.
Davies, Julian	Dept. of Computer Science, University of Western Ontario.
Dear, Patrick	"
Elcock, E.W.	"
Gascon, Jean	Dépt. de Psychologie, Université de Montréal.
Gileau, Francois	"
Godbout, G.	Dépt. d'Informatique, Université de Montréal.
Kasvand, Tony	Radio and Electrical Engineering, National Research Council, Ottawa, Ontario. K1A 0R8
Kuttner, Michael	Dept. of Computer Science, University of British Columbia, Vancouver, B.C., V6T, 1W5
Lemoyne, Giselle	Dépt. de Psychologie, Université de Montréal.
Mackworth, Alan	Dept. of Computer Science, University of British Columbia.

Mazlak, Lawrence J.

Dept. of Computing and Information
Science,
University of Guelph,
Guelph, Ontario, N1G 2W1

McCalla, Gordon

Dept. of Computer Science,
University of British Columbia.

Melli, Lucio F.

Dept. of Computer Science,
University of Toronto.

Montpetit, Guy

945 Chemin de Chambly,
Longueuil, Québec, J4H 3M6

Pylyshyn, Zenon

Dept. of Psychology,
University of Western Ontario.

Reich, Peter A.

Dept. of Computer Science,
University of Toronto.

Skuce, Douglas R.

Dept. of Electrical Engineering,
Université de Montréal.

Wong, Harry K.T.

Dept. of Computer Science,
University of Toronto.

Young, Douglas

Dept. of Computer Science,
University of Manitoba,
Winnipeg, Manitoba, R3T 2N2.

Zimmerman, Linda L.

Dept. of Computer Science,
University of Western Ontario.

