NEWSLETTER / MARCH 1985

CANADIAN ARTIFICIAL INTELLIGENCE 3

An official publication of CSCSI/SCEIO Canadian Society for Computational Studies of Intelligence/
Société canadienne pour l'étude de l'intelligence par ordinateur

The Symbolics 3600— a new computing paradigm.

Processing with symbols vs. computing with numbers. The first computers crunched numbers. And regardless of how powerful they've become, traditional systems still force you to deal with the world in quantitative terms.

Face sophisticated applications, and the limitations can become all too obvious.

An increasing number of computer scientists, researchers and program developers are discovering ways to break through this complexity barrier. Their vehicle—the Symbolics™ 3600.

The 3600 allows talented programmers and engineers to represent objects and knowledge far more flexibly than numeric formats allow.

Through the dynamic manipulation of arbitrary data structures consisting of symbols and their associated properties, a user can resolve problems previously considered impossible or infeasible.

SYMBOLIC PROCESSING WAS DEVELOPED TO SOLVE TWO KINDS OF PROBLEMS:

THE COMPLEX. AND THE IMPOSSIBLE.

A few typical applications. Custom VLSI engineering. The 3600 has a unique ability to deal with large, complex modeling structures. Semiconductor companies use it to assist in the development of the next generation of chips.

A development environment for complex software. Rapid prototyping and an incremental edit/compile/dynamic link/debug loop help make the 3600 one of the most effective programming environments ever developed.

Expert systems development. Using the 3600's powerful inferencing throughput and ability to handle very large knowledge bases, government agencies and Fortune 1000 companies are developing expert systems in such fields as process control, financial advisory services and image understanding.

Symbolics—the first name in symbolic processing. Symbolics was founded in 1980 to commercialize this new technology. Among the founders were the developers, at MIT, of the first hardware architecture designed specifically for symbolic processing.

Today, the 3600 represents the highest expression of symbolic processing technology. Its custom processor design incorporates a tagged memory architecture that manipulates symbols as efficiently as a traditional computer manipulates numbers.

Using Symbolics-enhanced Lisp, the 3600 provides a powerful integrated software environment which features object-oriented window system, incremental Lisp/Fortran-77/Pascal compilers and exceptional networking and graphics capabilities—for a new dimension in man-machine interaction.

To learn more about symbolic processing, write us at the address below.

symbolics inc.

Symbolics Inc., 11 Cambridge Center, Cambridge, MA 02142

NEWSLETTER
CANADIAN ARTIFICIAL INTELLIGENCE

An official publication of the Canadian Society for Computational Studies of Intelligence / Société canadienne pour l'étude de l'intelligence par ordinateur

NUMBER 3, MARCH 1985

Contents

Editor: Graeme Hirst
Assistance: Joanne Mager, Jim des Rivières
French/English translation: Jean-Pierre Corriveau, Yves Lespérance
Cover design: Cathy Ledden, University of Toronto Press
Cover logo: Kathryn Finter
Printing and Production Assistance: University of Toronto Press

Advertisers

Canadian Society for Computational Studies of Intelligence

Société canadienne pour l'étude de l'intelligence par ordinateur

CSCSI/SCEIO is the Canadian society for the promotion of interest and activity in Artificial Intelligence. It conducts workshops and fully refereed national conferences, publishes this newsletter, sponsors the journal *Computational Intelligence*, and coordinates activities with related societies, government, and industry.

To join CSCSI/SCEIO, use the membership form in this issue. Non-Canadian members are welcomed.

CSCSI/SCEIO is affiliated with the Canadian Information Processing Society and the International Joint Council on Artificial Intelligence.

CSCSI/SCEIO officers for 1984—86 are:

President:
Gordon McCalla
Department of Computational Science
University of Saskatchewan
Saskatoon, Sask S7N 0W0
Phone: 306-966-4902
UUCP: . . . !sask!kimnovax!mccalla

Vice-President:
John Tsotsos
Department of Computer Science
University of Toronto
Toronto, Ont M5S 1A4
Phone: 416-978-3619
UUCP: . . . !utcsri!utai!tsotsos

Secretary:
Michael Bauer
Department of Computer Science
University of Western Ontario
London, Ont N6A 3K7
Phone: 519-679-6048

Treasurer:
Wayne A. Davis
Department of Computing Science
University of Alberta
Edmonton, Alta T6G 2H1
Phone: 403-432-3976
UUCP: . . . !alberta!uqv-mts!sams

Canadian Artificial Intelligence Newsletter

Number 3 (New series) March 1985

ISSN 0823-9339

The *Canadian Artificial Intelligence Newsletter* is published quarterly by CSCSI/SCEIO, and is a benefit of membership in the society.

Editor:
Graeme Hirst
Department of Computer Science
University of Toronto
Toronto, CANADA M5S 1A4
Phone: 416-978-8747
CSNET: cscsi@toronto
ARPANET: cscsi.toronto@csnet-relay
UUCP: . . . !utcsri!cscsi
(UUCP connections to allegra, cornell, decvax, decwrl, deepthot, drea, floyd, garfield, hcr, ihnp4, linus, mbcsd, mcgill-vision, musocs, qucis, sask, ubc-vision, utzoo, uw-beaver, watmath, and many others.)

The *Newsletter* solicits contributions in English or French on any matter related to artificial intelligence, including:

- Articles of general interest.
- Descriptions of current research and courses.
- Reports of recent conferences and workshops.
- Announcements of forthcoming activities.
- Calls for papers.
- Book reviews (and books for review).
- Announcements of new AI companies and products.
- Opinions, counterpoints, polemic, controversy.
- Abstracts of recent publications, theses, and technical reports.
- Humour, cartoons, artwork.
- Advertisements (rates upon request).
- Anything else concerned with AI.

Please send submissions, either on paper or by network, to the editor or to your local *Newsletter* representative (see list on page 3). On-line submissions are preferred, but they should not contain justification spaces or hyphenated line breaks.

The *Newsletter* is published in March, June, September, and December. Material for publication is due on the 15th of the preceding month.

Please send changes of address to:
CSCSI/SCEIO, c/o CIPS
243 College Street, 5th floor
Toronto, CANADA M5T 2Y1

Editorial

Plus de français dans le *Canadian A. I. Newsletter*

Graeme Hirst
Editeur

Dans mes commentaires du dernier numéro, j'avais remarqué que "la domination anglophone au niveau de l'IA au Canada créait occasionnellement quelques problèmes embarassants." En fait, j'aurais dû écrire que "la domination anglophone au niveau de l'IA au Canada *est* un problème fort gênant".

Afin de rendre le *Canadian A. I. Newsletter* plus facile d'accès pour les francophones, nous avons decidé d'inclure dans chaque numéro, un résumé en français des principaux articles. Une traduction entière n'est guère pratique et l'anglais restera notre langue de travail. Nous espérons cependant que le résumé qui apparait à la page 7, aidera nos lecteurs francophones.

De plus, ce numéro comprend notre premier article en français (avec naturellement un résumé en anglais). Nous avons toujours encouragé la soumission d'articles dans l'une ou l'autre des deux langues, et c'est avec impatience que nous attendons un grand nombre de contributions des divers groupes francophones en IA au Canada.

□ □ □ □ □

En supplément spécial, *Vers un plan canadien de recherche en cinquième génération*, préparé par la Société canadienne de recherche en cinquième génération. Dans son article à la page 9, Eric Manning, président du comité qui a ébauché ce rapport, en explique l'origine et en justifie la présence dans ce bulletin.

Le rapport a été composé pour la société et ce bulletin avec l'aide de Stew Lee de l'université de Toronto, et la couverture a été conçue par Cathy Ledden de la University of Toronto Press. Nous les remercions de leur aide. □

L'échéance pour le numéro de juin est le 15 mai.

Editor's Notes

More French in *Canadian A. I. Newsletter*

Graeme Hirst
Editor

In my remarks in the previous issue, I said "The anglophone domination of AI in Canada occasionally leads to embarrassing problems". What I should have said was "the anglophone domination of AI in Canada *is* an embarrassing problem".

In order to make the *Canadian A. I. Newsletter* more accessible to francophones, we are now including in each issue a summary in French of the main articles. A complete translation of the issue is not practical, and English will remain our working language, but we hope the summary will assist our francophone readers. It appears on page 7.

In addition, this issue carries our first article in French (with, of course, an English summary). We have always solicited articles in either language, and look forward to more contributions from francophone AI groups in Canada.

□ □ □ □ □

Enclosed with this issue is a special supplement, *Towards a Canadian Fifth Generation Research Plan*, produced by the Canadian Society for Fifth Generation Research. In an article on page 9, Eric Manning, chairman of the committee that drafted the report, explains how it came to be and why it was thought appropriate to distribute it as a supplement to the *Newsletter.*

The report was typeset for the Society and the *Newsletter* with the assistance of Stew Lee of the University of Toronto, and the cover was designed by Cathy Ledden of the University of Toronto Press. We are grateful for their assistance. □

Deadline for the June issue is 15 May.

Readers Write

Letters to the Editor

Catechisms

Some readers may appreciate a brief background of *catechisms*, the subject of Alan Mackworth's parody ("A Catechism for the Neat AI Person", *Canadian Artificial Intelligence Newsletter*, December 1984).

Catechisms are manuals for religious instruction. They have been published since the Middle Ages; authors of catechisms include Maimonides and Luther. While the contents of different catechisms vary, they typically proclaim faith in the existence of God, and include the Ten Commandments.

While there is quite a difference between historical catechisms and Dr Mackworth's, it is worth noting that faith is an important element in our lives and in the pursuit of scientific study.

Brian Nixon
University of Toronto

Alan Mackworth sent the following "exegesis" in response to the preceding letter:

The context for my catechism can best be captured by recalling the hoary old drunk-and-the-keys joke. Skip the next paragraph if you've heard it.

John McCordy was wandering home from the pub one dark night when he dropped the key to his house. He kept going until he got to a streetlight and started to crawl around on his hands and knees looking for it. When a cop stopped to ask him why he didn't look where he had lost the key, John replied, "Don't be silly, it's pitch bloody black over there!"

This story illuminates the fundamental tension in any science between theory and reality. That gap produces the underlying dialectic that drives the scientific process. (Of course, the story is usually told by experimentalists to ridicule theoreticians who promote sterile formalisms with unrealistic restrictions.)

The theory—reality gap in AI is a chasm. Actually, it is many chasms. In knowledge representation there are plenty of streetlights, but they are all far removed from the key. It is important for the drunks under any given streetlight to keep up their spirits as they grovel around. This they do by singing hymns and chanting catechisms. Given the political and sociological nature of scientific activity it is also important that new drunks coming onto the scene be convinced that there is only one true streetlight and that its wattage is increasing. Other streetlights are either pale reflections of ours or total mirages. The declarative—procedural controversy is a classic case study in behaviour of this kind. There is only one way to enter God's kingdom.

So my harmless little catechism was a not-very-funny parody designed to remind us that AI paradigms are like religious sects that must keep up the faith even when faced by overwhelming isolation and rejection or, as is more common these days, tempted by Mammon.

Alan Mackworth
University of British Columbia

Kind words for the *Newsletter*

The current *Canadian A. I. Newsletter* looks great! The editor's note about the name change was pretty funny. I'm glad you coerced me into subscribing.

John Martin
ITT Advanced Technology Center
Shelton, Conn., U. S. A.

Newsletter's electronic address changes

The UUCP mail node previously known as utcsrgv has changed its name to utcsri. The *Newsletter*'s UUCP address therefore becomes . . . !utcsri!cscsi. The old address will continue to work for a short time. Other networks are not affected.

Guidelines for electronic submissions

The *Canadian A. I. Newsletter* prefers submissions to be made by network mail when possible, to eliminate the need for retyping. However, please do not send text that has been through a document formatter; extra embedded spaces and hyphenated line breaks just have to edited out again before typesetting. The best thing to send is plain typing or *troff* source (preferably with —me macros).

Canadian Artificial Intelligence Newsletter

Numéro 3, mars 1985

Résumé

Annonceurs

YOUR GATEWAY TO ARTIFICIAL INTELLIGENCE . . .

GOLDEN COMMON LISP

- Golden Common Lisp (**GCLisp**) is the first implementation of industry standard Common Lisp for personal computers such as IBM PC, XT, and AT computers and their compatibles

- The $675 (Canadian) **GCLisp** package includes the following:

  - the **GCLisp interpreter**
  - the **GMACS** intelligent editor
  - a Common Lisp **tutorial subsystem**
  - an **on-line help** of all the Lisp and editor functions
  - a standard **Common Lisp Manual** by Guy Steele
  - the second edition of P. Winston's text book, LISP
  - the **GCLisp** Manual

- System requirements:

  - minimum 512 KByte memory
  - 100% compatibility with the IBM PC, XT, or AT
  - PC-DOS or MS-DOS version 2.0 or later

- Of over 20 Lisp dialects, Common Lisp is now supported by most major AI R&D groups in the U.S. They include: Stanford, Carnegie-Mellon, and Yale Universities, MIT, U.C. Berkley, Digital Equipment Corporation, Bell Laboratories, Lawrence Livermore Laboratories, Symbolics Inc., Lisp Machines Inc., Xerox Corporation, Hewlett-Packard Corporation, Texas Instruments, and Perq Systems Corporation. The U.S. government (DARPA) also backs Common Lisp.

- The tutorial subsystem introduces the user to powerful Lisp programming techniques and guides him/her through the learning process. **GCLisp** package is more than a language system. It is an affordable AI Development Environment.

- **GCLisp** facilitates rapid prototyping of pilot AI systems such as rule-based Expert Systems, Natural Language Processing systems, and other Intelligent Interface systems.

- **GCLisp** package is available from Applied AI Systems Inc. for $675 Canadian (PST not included). Shipping charges will be paid by Applied AI Systems Inc., if payment accompanies the order. Optional technical support is available from Applied AI Systems Inc. (Price may change without notice)

Applied AI Systems Inc. is an authorized dealer of products by Gold Hill Computers, Inc. of Cambridge Mass. **GCLisp** is a registered trademark of Gold Hill Computers, Inc. IBM PC, XT, AT, and PC-DOS are registered trademarks of International Business Machines Corporation. MS-DOS is a registered trademark of Microsoft, Inc.

AI
Applied AI Systems, Inc.
P.O. Box 13550
KANATA, Ontario
Canada K2K 1X6
Telephone: (613) 592-0084

The Canadian Society for Fifth Generation Research, and Its Plan

Eric Manning
Chairman, Steering Committee
Canadian Society for Fifth Generation Research

In 1982 the Japanese Ministry of International Trade and Industry (MITI) produced its concept and program of Fifth Generation Computing. The program was built upon microelectronics, computer architecture, computer systems and languages (including operating systems and databases), and artificial intelligence. Its basic premise was that these topics have matured to a point that permits the development of knowledge-based systems, featuring logical inference, communication with humans through natural language, and so forth. The task therefore is essentially one of systems integration, and the budget allocated was about $500 million over 10 years.

The reaction from the West was rapid and dramatic; the British committed about $750 million over three years to a program of computing research based on microelectronics, architecture, systems and languages, and AI (the "Alvey Programme"); the EEC committed 1.2 billion ECU to a similar program, called ESPRIT; the U.S. began both its Microelectronics and Computer Technology Corp (MCC) and its Strategic Computing Initiative ($1 billion); West Germany started a program and the Australians have proposals in the works. In Canada, there was only silence.

Deeply concerned that so-called Fifth Generation computing might turn out to be a key development, and that nations which failed to participate in its development might lose expertise crucial to the future development of manufacturing, education, resource exploitation, etc., a small group of Canadian academics asked the Natural Sciences and Engineering Research Council (NSERC) to sponsor a meeting. NSERC did so, and about 100 academics representing all regions of Canada met in Ottawa last March. They were joined by about 30 representatives of Canadian industry and government; all present agreed to constitute themselves as the Canadian Society for Fifth Generation Research (CSFGR), and a Steering Committee was elected and asked to prepare a draft plan for the academic component of a Canadian research program.

The Committee sent out a call for draft research proposals; it received about thirty draft proposals, which were edited into a draft plan and sent to the membership for approval or otherwise. The third draft was overwhelmingly approved by the members, and therefore became the CSFGR's Plan. It is reproduced as a supplement to this issue of the *Canadian A.I. Newsletter*, in order that it be widely distributed and discussed.

Unfortunately, the Plan is in reality more a compendium than a plan; it lacks focus and cohesion. Also, the support of Canadian industry is absolutely essential if the Plan is to receive government funding. (Any funding received would be dispensed through the peer-review process of NSERC.) To address both of these problems, members of the Steering Committee are approaching Canadian firms, inviting them to provide support (money, people, equipment, etc.) to the projects of their choice. The President of NSERC supports this course of action and has asked to be kept informed of its results.

The CSFGR is currently working with NSERC to obtain a travel subsidy for the Society's next meeting (about $20,000 is required), and operating funds for the Society. An Executive needs to be elected and some kind of Constitution or formal statement of purpose needs to be drafted and discussed. Finally, the problem of cohesion will require continuing attention. Perhaps some kind of mechanism to periodically canvass the views of both academic and industrial members, and from these produce draft recommendations for NSERC concerning topics particularly deserving of extra support, can be devised.

People interested in being added to the Society's mailing list should write to the author at Institute for Computer Research, University of Waterloo, Waterloo, Ont N2L 3G1. □

Theoretical Issues in Natural Language Understanding

There has been an increasing interest in natural language processing over the past few years, due in part to the tremendous resurgence of interest in artificial intelligence research, particularly natural language understanding research, and the importance of natural language understanding for Fifth Generation computing projects. It has been some number of years since the Theoretical Issues in Natural Language Processing (TINLAP) workshops have been held. CSCSI/SCEIO decided that it was time to do something about this state of affairs, and accordingly has organized a workshop to be held at Dalhousie University, Halifax, Nova Scotia, 28—30 May 1985.

Theoretical Approaches to Natural Language Understanding (TANLU) is to bring together active researchers in Computational Linguistics, Artificial Intelligence, Linguistics, Philosophy, and Cognitive Science to hear and discuss invited talks, papers, and positions relating to some of the current "hot" issues in natural language understanding. Three topics will form the focus for discussion: aspects of grammars, aspects of semantics, and knowledge representation.

Each of these topics will consider current methodologies:

In grammars:
theoretical developments, especially generalised phrase structure grammars and logic-based meta-grammars;

In semantics:
situation semantics and Montague semantics;

In knowledge representation:
logical systems (temporal logics, etc.) and special purpose inference systems.

A day will be devoted to each of the three; a preliminary program appears on the following pages. The organizers of each session have given the invited speakers a number of questions around which to focus their talks, and these are printed with the program.

Forms to register for the workshop and reserve accommodation are enclosed with this *Newsletter.* A discount for early registration will apply until 15 April; dormitory accommodation must be reserved by 7 May. If the forms are missing, or of you require more information, contact:

Richard Rosenberg
Dept Mathematics and Computing Science
Dalhousie University
Halifax, NS, CANADA B3H 4H8
Phone: 902-424-2572

In the event of a Canadian postal strike, you may pre-register for the workshop and reserve accommodation by phone or electronic mail, and pay upon arrival.

UUCP: . . . ! {utcsri, dartvax}! dalcs! rsr
CSNET: cscsi@toronto
ARPA : cscsi%toronto@csnet-relay

Following are the organizers of TANLU:

General Chairperson:
Richard Rosenberg, Dalhousie University

Program Chairperson:
Nick Cercone, Simon Fraser University

Local Arrangements:
Jan Mulder, Dalhousie University

Program/Organising Committee:
Len Schubert, University of Alberta
Veronica Dahl, Simon Fraser University
David Israel, BBN Labs & SRI International
Graeme Hirst, University of Toronto
Ralph Weischedel, BBN Laboratories
James Allen, University of Rochester

Invited Speakers and Panelists:
Harvey Abramson, University of British Columbia
Robin Cooper, University of Wisconsin
Dan Flickinger, Hewlett-Packard
Pat Hayes, University of Rochester
Don Hindle, Bell Labs
Lynette Hirshman, System Development Corporation
Ron Kaplan, Xerox PARC
Mitch Marcus, Bell Labs
Bill Mark, Savoir
Eric Mays, IBM Thomas J. Watson Research Center
Fernando Pereira, SRI International
Stan Peters, CSLI, Stanford University
Stan Rosenschein, SRI International
Paul Sabatier, Paris
Patrick Saint Dizier, IRISA, Beaulieu
Candy Sidner, BBN Laboratories
Norm Sondheimer, USC / ISI
David Scott Warren, SUNY Stony Brook
Dave Touretzky, Carnegie-Mellon University
William Woods, Applied Expert Systems

Theoretical Issues In Natural Language Understanding

Dalhousie University, Halifax, Nova Scotia
28–30 May 1985

Preliminary Schedule

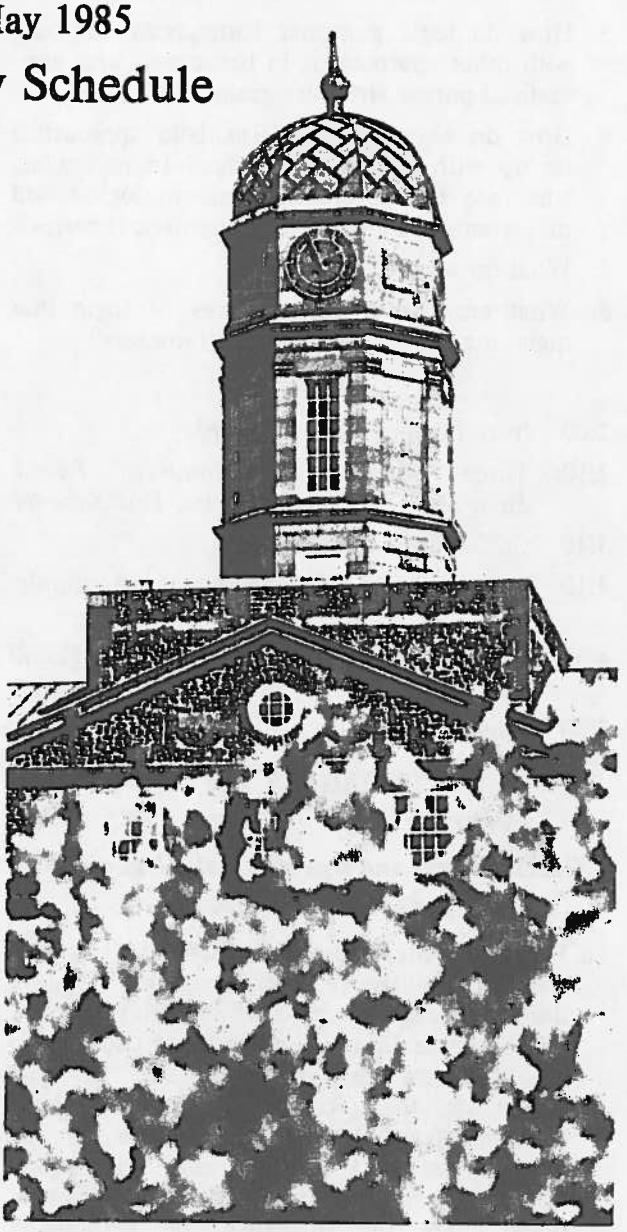

Dalhousie University

Grammar Day
Tuesday 28 May 1985

Morning

Generalised Phrase Structure Grammars
Len Schubert, Organiser

Questions for discussion

1. What are the desiderata in characterizing *(a)* natural languages, as abstract objects; *(b)* natural language understanding and generation, as abstract processes (*i.e.*, viewed in a more or less machine-independent way)?
2. What's good about each grammatical framework (Generalised Phrase Structure Grammars, Government–Binding, Lexical Functional Grammars) with respect to 1*a*?
3. What's good about each parsing or generation framework with respect to 1*b*?
4. What are the most important areas for further work on 1*a*?
5. What are the most important areas for further work on 1*b*?

Schedule

9:00 Introduction, *Len Schubert*

9:10 Two invited presentations of 40 minutes each, *Ron Kaplan, Dan Flickinger*

10:30 Coffee break

10:50 Invited presentation, *Mitch Marcus and Don Hindle*

11:30 Two presentations of 25 minutes each, *TBA*

12:30 Luncheon

Afternoon

Logic-Based Meta-Grammars
Veronica Dahl, Organiser

Questions for discussion

1. Are logic grammar formalisms convenient or necessary? Would plain Prolog be a better alternative?
2. In view of the justification of grammar formalism, what features would you absolutely retain from each in order to develop a single, comprehensive one? Why?
3. What new features is it desirable to provide logic-based metagrammars with? Why? How feasible are they?
4. What are the inherent limitations of logic-based metagrammars for processing language?

Theoretical Issues In Natural Language Understanding

Preliminary Schedule (continued)

5. How do logic grammar formalisms compare with other approaches, in particular, with generalised phrase structure grammars?
6. How do computational linguistic approaches tie up with linguistic theories? In particular, what are the common points in logic-based metagrammars and current linguistic theories?
7. What do we mean by logic?
8. What are the critical features of logic that make it attractive for writing grammars?

Schedule

2:00 Introduction, *Veronica Dahl*

2:10 Three 30-minute presentations, *Patrick Saint-Dizier, Fernando Pereira, Paul Sabatier*

3:40 Coffee break

4:10 A paper to be chosen from those submitted.

4:30 Panel Discussion, *Harvey Abramson, Veronica Dahl, Lynette Hirschman, Paul Sabatier*

5:30 Finish

Semantics Day
Wednesday 29 May 1985

Graeme Hirst and David Israel (Organisers)

Questions for discussion

1*a*. What does Montague semantics have to say to artificial intelligence, and to natural language understanding in particular? Which aspects of Montague semantics are good and useful, and which aspects are bad and useless? To what extent can the good parts be used without problems from the bad parts? Can the bad parts be salvaged by appropriate adaptation? For what particular natural language applications might the formalism be particularly appropriate or particularly inappropriate?

1*b*. Ditto for situation semantics.

2. By an AI representation is meant a semantic representation or formalism that permits structuring of knowledge (IS-A, slots, scripts, schemata, and all that), incremental storage of facts, retrieval, inference, abductive reasoning, etc. What is the relationship between the two semantic formalisms and AI representations such as slots, frames, or production rules? To what extent could each of the formalisms *be* such a representation? That is, could they be extended to permit a suitable structuring, inference and abduction mechanisms, update and retrieval, and do it all with adequate efficiency? What is necessarily lost? (This is especially critical for situation semantics, for which 'computability' claims are made but of which, unlike the systems of higher-order logic upon which Montague semantics is based, little study of the AI properties has yet occurred.)

   Consider, for example, an expert system for medical diagnosis that keeps itself up-to-date by reading the *New England Journal of Medicine* each week. There will need to be some commonality between the semantic formalism for the natural language input and the rules (or other representation) used for medical diagnosis. Can we see how either of the formalisms could be used in this? If an intermediary process to translate or interpret the formalism into rules must be posited, then what has been gained by the use of the formalism?

   In a different domain: Could either of the formalisms be (the basis of) a suitable interlingua for a machine translation system? Are they suitable for generating natural language from? What aspects of the input that must necessarily be preserved in the output, such as discourse focus and tenor, are not dealt with by the formalisms?

   Similar questions also arise in most other applications of NLU systems. To what extent could the formalisms aid the solution of traditional NLU problems such as reference, ambiguity, and determining discourse connections? (*Cf* Hirst's Montague-based semantic interpreter, in which ambiguity resolution, both lexical and structural, had to be transparent to the Montague parts of the system.)

3. What alternative formalisms are available or foreseeable?

Schedule

9:00 Neutral, historical overview of, and introduction to, semantic theories, especially as applied to natural language understanding, *David Israel and Graeme Hirst*

9:40 Why I like Montague semantics, and what they have for natural language

Theoretical Issues In Natural Language Understanding

Preliminary Schedule (continued)

understanding, *TBA*

10:50 Coffee break

11:20 Why I like situation semantics, and what they have for natural language understanding, *Stan Peters*

12:30 Luncheon

2:00 Three half-hour responses to the preceding, *William Woods, Norm Sondheimer, Robin Cooper*

3:30 Coffee break

4:00 Another response, *David Scott Warren*

4:20 Discussion time, starting with re-responses from morning speakers

5:30 Finish

Knowledge Representation Day
Thursday 30 May 1985

Ralph Weischedel and James Allen (Organisers)

Questions for discussion

1. What are the general-purpose and special-purpose approaches? How and why do they differ?

   In the past, one could try to pigeonhole and stereotype general-purpose approaches as those concerned with formal analysis, involving axiomatisation and also characterisation of the class of legal inferences given a base set of facts. Concern for run-time computational issues was treated as separate from issues of expressive power. A theory could in principle be run using a theorem prover (*i.e.*, an actual theorem prover, or a general parsing algorithm). Special-purpose approaches were concerned with describing a mechanism that performs a certain set of inferences in a given setting. The focus had been on the specific mechanisms needed for the sake of "computational efficiency", cognitive modelling, or simplicity (in the sense of not using a cannon to kill a fly). Formal concerns, such as expressive adequacy, semantics of the formalism, and characterisation of the class of possible inferences were of less concern.

   Recent work has blurred this distinction, as specific mechanisms have been analysed as formal theories, *e.g.*, Touretzky's analysis of NETL inheritance, Brachman and Levesque's study of tractability of inheritance, Frisch's formalisation of retrieval as limited inference, and so on.

2. Select at least one of the following and defend your belief regarding which of general theorem provers and special-purpose inference engines are "better" for the phenomenon:
   - lexical ambiguity (*e.g.*, recognising what sense of "saw" or "of" is being used in a context).
   - modifier attachment (*e.g.*, determining the attachments of the PPs (prepositional phrases) in a structure like V NP PP PP).
   - anaphora (*e.g.*, determining what is meant by an anaphoric expression).
   - contextual ellipsis (*e.g.*, identifying the meaning of both the input and elided material).
   - figures of speech (*e.g.*, what is meant in metonymy, an analogy, etc.).
   - intention (*e.g.*, identifying the intention of an utterance by a speaker in context).
3. What aspects of common sense reasoning argue for general or special inference engines?
4. How can we have real-time language processing using general or special inference engines?
5. What aspects of reasoning (if any) for language processing are not covered by common sense reasoning?

Schedule

9:00 Overview, *James Allen and Ralph Weischedel*

9:30 Invited Speaker, *Pat Hayes*

10:20 Coffee break

10:50 Invited Speakers, *Stan Rosenschein, Eric Mays*

12:30 Luncheon

2:00 Invited Speakers, *Bill Mark, Candy Sidner*

3:40 Coffee break

4:10 Invited Speaker, *Dave Touretzky*

4:50 Discussion time

5:30 Finish

LISP Canada Inc.

Distributes and Services the LMI Lambda Machine

LMI Lambda Machine

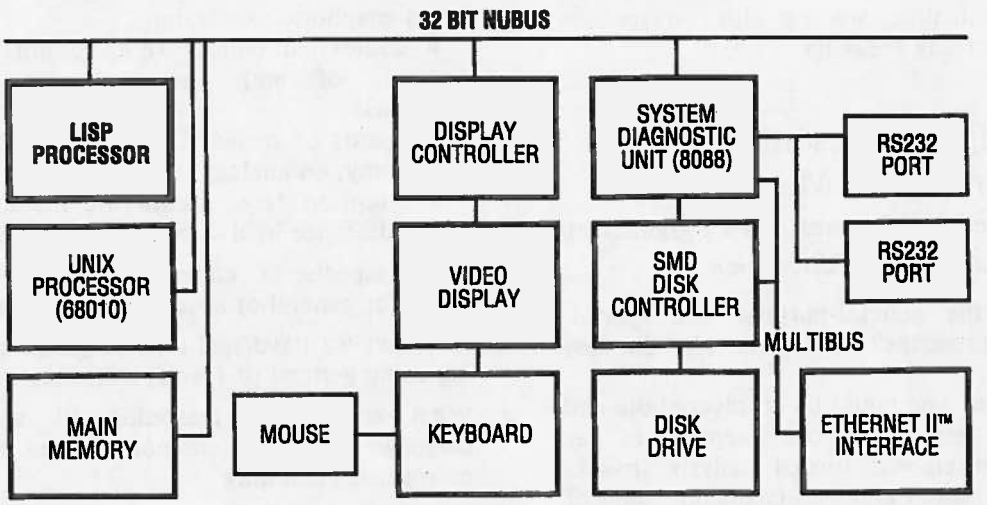

USERS BENEFIT FROM:

- MULTIBUS™ Compatibility
- Multi-processor Capability
  - LISP Processor
  - 68010 UNIX Processor
  - 8088 System Diagnostic Unit
- Virtual Control Store and Microcompiler
- Powerful Software Resources

LISP Canada Inc.
5252 de Maisonneuve W. Blvd.
Montreal, Quebec, Canada
H4A 3S5
Tel: 514 **487 7063**
Telex: 055-66382

LISP Canada inc.

Distribue et entretient la machine Lambda de LMI

Machine Lambda de LMI

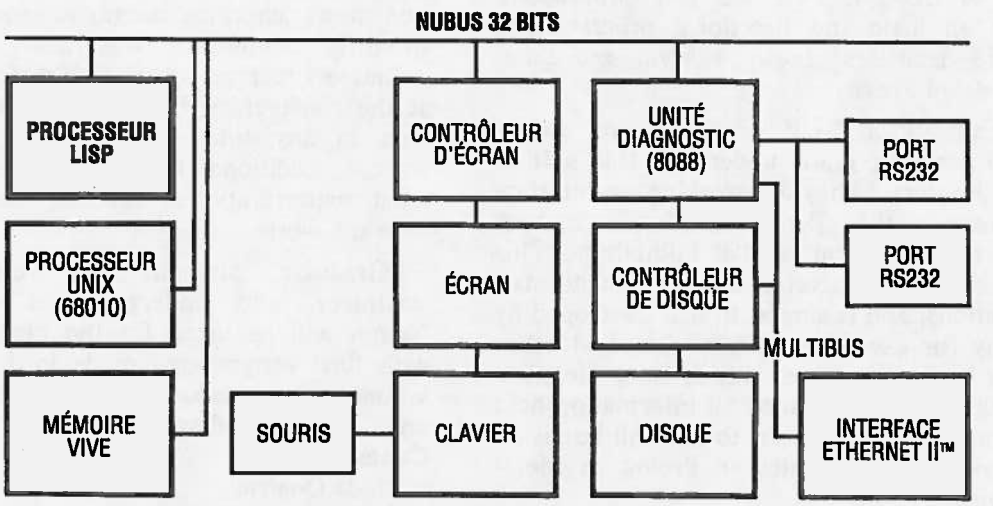

LES AVANTAGES DU LAMBDA

- **Compatibilité MULTIBUS™**
- **Accès à de multiples processeurs**
  - **Processeur LISP**
  - **Processeur UNIX (68010)**
  - **Unité de diagnostic (8088)**
- **Microcompilateur**
- **Logiciels puissants**

LISP Canada Inc.
5252 de Maisonneuve W. Blvd.
Montreal, Quebec, Canada
H4A 3S5

Tel: 514 **487 7063**
Telex: 055-66382

Visitor at University of Western Ontario

Mike Bauer

Prof Peter M.D. Gray from the University of Aberdeen in Scotland is visiting at the University of Western Ontario on sabbatical until June, at the invitation of Prof Ted Elcock, chairman of the Computer Science Department. Prof Gray worked with Prof Elcock in Aberdeen from 1968 to 1971 on the ABSYS and ABSET systems, which were forerunners of Prolog. He implemented a database on disk for general list structures with named components. Subsequently he has worked on automatic program generation for access to Codasyl databases, implemented using Prolog, as reported at IJCAI-83. He has just published a textbook on logic and functional programming applied to databases: *Logic, Algebra, and Databases* (Halsted Press).

Currently, Prof Gray heads a team with a $200,000 research grant under the U.K.'s IKBS (Alvey) program. They are working on interfacing Prolog to the 'Persistent Heap' database developed by Atkinson *et al* at Edinburgh. This is suitable for AI structures and not just normal-form relations, and is similar to that developed by Prof Gray for use with ABSYS. Amongst other things, it is planned to use this to store Metadata describing scientific and medical information held in remote databases. Access to this will be via an expert system implemented in Prolog: A Metadata Advisor.

Prof Gray would be very interested to make contacts with groups with similar interests and is happy to give invited talks.

Phone: 519-679-2976
UUCP: . . .!watmath!deepthot!pmdg or
. . .!utcsrgv!utzoo!uwo!pgray

IJCAI seeks student workers

The Ninth International Joint Conference on Artificial Intelligence will be held from 18 to 24 August 1985, at the University of California at Los Angeles.

The IJCAI Local Arrangements Committee is looking for student volunteers for the conference. Volunteers work approximately 8 hours during the conference. Tasks include manning information desks, checking badges at sessions, and distributing conference materials. In exchange volunteers receive a staff T-shirt, free registration at the conference, Proceedings, and free admission to any tutorials at which the volunteer works. Additional benefits include a party and great opportunity for meeting people from all over the world.

Graduate students are encouraged to volunteer, and undergraduates are welcome. Names will be taken for the next few months, with final assignments made in July. Tentative volunteers are welcome. Volunteers will be taken on a first-come/first-served basis, so reply now. Contact:

Linda Quarrie
The Robotics Institute
Carnegie-Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213, U.S.A.
Phone: 412-578-8815 or 412-521-1968
ARPANET: lindaq@cmu-ri-isl1.arpa

Advertising in the *Canadian A. I. Newsletter*

An advertisement in the *Canadian A. I. Newsletter* reaches CSCSI/SCEIO's 535 (and rapidly growing) members in the AI R&D community, and also key people in the Canadian media and federal and provincial governments.

Our rates are extremely reasonable; $^{CDN}100 ($^{US}75) buys a full-page black-and-white advertisement. For full details, or to book space, contact the editor, Graeme Hirst, 416-978-8747.

Cartoon by Johnny Ng, University of Alberta

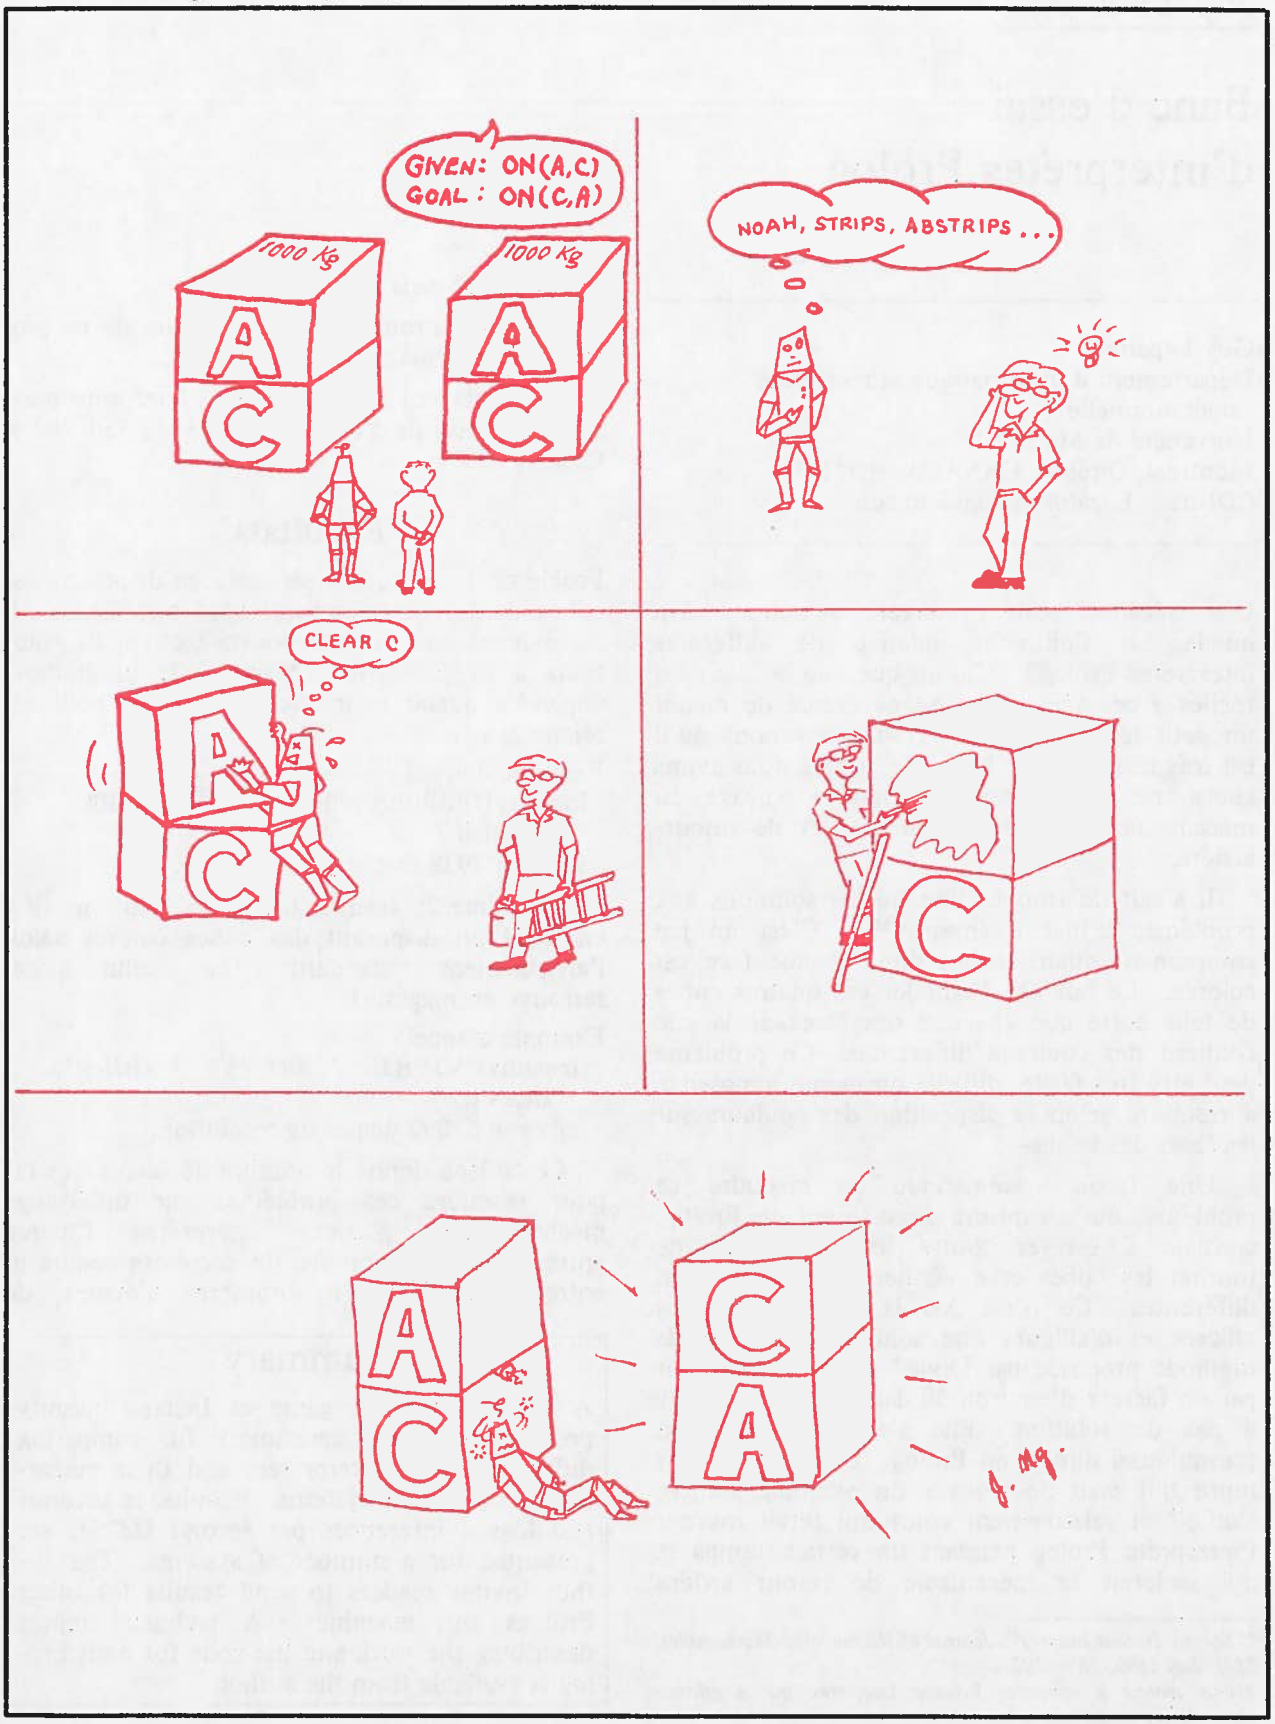

Banc d'essai d'interprètes Prolog

Guy Lapalme
Département d'informatique et recherche opérationnelle
Université de Montréal
Montréal, Québec, CANADA H3C 3J7
CDNnet: Lapalme%iro.udem.cdn

Une question revient souvent a notre esprit: quelle est l'efficacité relative de différents interprètes Prolog? N'ayant que peu de données réelles à cet effet, nous avons décidé de mener un petit test comparatif. Nous convenons qu'il est très restreint, car le problème que nous avons choisi ne vérifie en fait que la vitesse du mécanisme d'unification (simple) et de retour-arrière.

Il s'agit de trouver une ou les solutions aux problèmes d'Instant Insanity[MD]. C'est un jeu comprenant quatre cubes dont chaque face est colorée. Le but est d'empiler ces quatres cubes de telle sorte que chacune des faces de la pile contient des couleurs différentes. Ce problème peut être très facile, difficile ou même impossible à résoudre selon la disposition des couleurs sur les faces des cubes.

Une façon systématique de résoudre ce problème, qui s'exprime directement en Prolog, consiste à essayer toutes les possibilités de tourner les cubes et à vérifier si les faces sont différentes. Ce n'est pas la solution la plus efficace et d'ailleurs une solution basée sur la méthode proposée par Levin* permet d'accélérer par un facteur d'environ 50 dans un cas où il n'y a pas de solution. Elle s'exprime de façon presqu'aussi directe en Prolog. De toutes façons notre but était de trouver un exemple simple, "utile" et relativement court qui ferait tourner l'interprète Prolog pendant un certain temps et qui isolerait le mécanisme de retour arrière fondamental dans Prolog.

Nous montrons donc une version de ce programme en syntaxe "style DEC-10".

Les tests ont été faits avec l'interprète portable en Pascal de York sur VAX/VMS 750/780 et CYBER/835.

Resultats†

Problème 1: tenter de résoudre en disposant des cubes dont toutes les faces sont semblables. Il est évident que l'on ne pourra trouver de solutions à ce problème. Mais ici le programme engendre quand même les possibilités pour les rejeter aussitôt.

Exemple d'appel:
insanity("BBBBBB", "BBBBBB", "BBBBBB", "BBBBBB")
environ 7938 étapes de résolution.

Problème 2: trouver toutes les solutions (il y en a 2) en disposant des cubes colorés selon l'arrangement "standard" (*i.e.*, celui qu'on retrouve en magasin).

Exemple d'appel:
insanity("VJVRJB", "JBRBVV", "RVBJRR", "BRJVRJ").
environ 52092 étapes de résolution.

Le tableau donne le nombre de secondes CPU pour résoudre ces problèmes sur différentes machines avec plusieurs interprètes. Chaque entrée donne le nombre de secondes requis et entre parenthèses le nombre d'étapes de

Summary

A program for the game of Instant Insanity provides a simple benchmark for comparing different Prolog interpreters and their performance on various systems. Results, in seconds and logical inferences per second (LIPS), are presented for a number of systems. The author invites readers to send results for other Prologs and machines. A technical report describing the work and the code for each Prolog is available from the author.

*"*Solving Instant Insanity", Journal of Recreational Mathematics, 2(3), July 1969, 189–192.*

†Nous tenons à remercier Roxane Lacouture qui a effectué plusieurs des essais relatés ici.

Machine	Prolog	Prob 1		Prob 2	
		Sec	(LIPS)	Sec	(LIPS)
CYBER-835	IIUW	18.5	(429)	98.1	(531)
	York	130.0	(61)	493.0	(105)
VAX/750	Prolog-II	116.2	(68)	570.5	(91)
	York	223.4	(35)	701.5	(74)
VAX/780	Prolog-II	59.4	(133)	297.6	(175)
	York	142.0	(55)	450.0	(115)
	M-Prolog	34.0	(233)	197.0	(264)
Osborne-I	Micro-Prolog	120.0	(66)	730.0	(71)
IBM PC	Micro-Prolog	154.0	(51)	916.0	(56)

résolution par seconde; cette dernière valeur est équivalente aux LIPS (logical inferences per second).

Ce tableau permet de dégager les conclusions (grossières) suivantes. Pour un même problème et un même interprète:

- La vitesse de résolution sur un VAX/750 est environ la moitié de celle sur VAX/780.
- La vitesse de résolution sur le CYBER-835 est la même que sur le VAX-780.
- La vitesse de résolution de Micro-Prolog sur Osborne-I est environ 20% plus rapide que sur IBM PC.

Entre les interprètes:

- Prolog-II est deux fois plus rapide que York et M-Prolog prend environ 2/3 du temps de Prolog-II.
- Le temps réel de Micro-Prolog sur Osborne-I équivaut à peu près au temps CPU de Prolog-II sur VAX/750 ou York sur VAX/780. Ce point est très suprenant!

Nous ne demandons pas mieux que de compléter ce tableau avec d'autres machines et d'autres interprètes. Si vous avez accès à de tels outils, nous vous suggérons fortement d'adapter le programme à votre syntaxe tout en conservant le même ordre de clauses dans les "procédures" et en respectant la structuration, même si votre Prolog préféré possède des "facilités" qui pourraient accélérer quelque peu le temps de solution.

Nous apprécierions beaucoup que vous nous communiquiez vos résultats. Le texte complet de ces tests ainsi que de quelques autres constitue le document de travail #151 du Département IRO, Université de Montréal. □

```
/* Banc d'essai Prolog: Trouver les solutions au jeu
   d'Instant Insanity de la façon brutale, i.e., générer
   toutes les solutions. */

conc2([A,B],[C,D],[A,B,C,D]).

selection1([F1,F2,F3,F4,F5,F6],[F1,F2],[F3,F4,F5,F6]).
selection1([F1,F2,F3,F4,F5,F6],[F3,F4],[F1,F2,F5,F6]).
selection1([F1,F2,F3,F4,F5,F6],[F5,F6],[F1,F2,F3,F4]).

selection2([FA,FB,FC,FD],[FA,FB]).
selection2([FA,FB,FC,FD],[FC,FD]).

rotation(TOUT,TOUT).
rotation([DEV,DER,HAU,BAS],
  [HAU,BAS,DER,DEV]).
rotation([DEV,DER,HAU,BAS],
  [DER,DEV,BAS,HAU]).
rotation([DEV,DER,HAU,BAS],
  [BAS,HAU,DEV,DER]).

disj([],[]).
disj([T1|Q1],[T2|Q2]):- T1=\=T2, disj(Q1,Q2).

verifier(SOL1,SOL2,SOL3,SOL4) :-
  rotation(SOL2,C2), disj(SOL1,C2),
  rotation(SOL3,C3), disj(SOL1,C3), disj(C2,C3),
  rotation(SOL4,C4), disj(SOL1,C4), disj(C2,C4),
    disj(C3,C4),
  sorm("cube 1 : "),sorm(SOL1),nl,
  sorm("cube 2 : "),sorm(C2),nl,
  sorm("cube 3 : "),sorm(C3),nl,
  sorm("cube 4 : "),sorm(C4),nl.

insanity(C1,C2,C3,C4) :-
  selection1(C1,P1A,R1),
  selection1(C2,P2A,R2),
  selection1(C3,P3A,R3),
  selection1(C4,P4A,R4),
  selection2(R1,P1B),
  selection2(R2,P2B),
  selection2(R3,P3B),
  selection2(R4,P4B),
  conc2(P1A,P1B,SOL1),
  conc2(P2A,P2B,SOL2),
  conc2(P3A,P3B,SOL3),
  conc2(P4A,P4B,SOL4),
  verifier(SOL1,SOL2,SOL3,SOL4).

sorm([]).
sorm([T|Q]) :- put(T),sorm(Q).
```

New journal

First issue of *Computational Intelligence*

By the time this issue of the *Canadian A.I. Newsletter* is in print, the Spring issue of Volume 1 of *Computational Intelligence* will have appeared. This is the first issue of the National Research Council of Canada's 13th journal, and it is also the first new journal supported by the NRCC in twelve years. The journal is sponsored by CSCSI/SCEIO, and edited by Nick Cercone and Gordon McCalla.

Computational Intelligence aims to improve communication between the many specialised fields of research in artificial intelligence and thus stress the cohesiveness of AI as a field of research. The journal also aims to reach a wide audience. In a departure from the traditions of NRCC, which stresses the Canadian origins of its twelve other journals, *Computational Intelligence* is advertised as an international journal. The international character of the journal is reflected in the composition of its editorial board.

Computational Intelligence has a general interest in AI, and it will publish high-quality original theoretical or experimental research papers in artificial intelligence. The editors also want to encourage submissions that report on experimental results, and on philosophical speculations in AI. It is the aim of *Computational Intelligence* to make a wide range of research and theoretical investigations available and accessible to a general artificial intelligence audience.

The first issue reflects the wide scope of the journal: it contains papers that discuss natural language problems in AI, knowledge representation, and reports on applications of expert systems. The four papers presented in the first issue have been contributed by Candace L. Sidner; John K. Tsotsos; David Etherington, Robert E. Mercer, and Raymond Reiter; and David Wilkins. Abstracts are given below. The first issue will also explain the processes through which papers pass from submission to publication, and will include information for potential authors.

Special issues will follow each year; work is in progress on these issues already, and editorial responsibilities have been assigned. The first special issue will be devoted to the Theoretical Approaches to Natural Language Understanding based on the workshop sponsored by CSCSI/SCEIO in Halifax during May 1985, and will be edited by Nick Cercone. AI applications in education, edited by Dr. Marlene Colbourn of the University of Waterloo will form a second special issue. The third will be machine learning, edited by Dr. Larry Rendell of the University of Illinois. The tentative schedule for these special issues is as follows:

Natural language, 1(4), 1985
AI in education, 2(3), 1986
Machine learning, 3(1), 1987

To subscribe to *Computational Intelligence*, use the subscription form included with this issue of *Canadian A.I. Newsletter.* Members of CSCSI/SCEIO receive a considerable discount on the regular subscription price.

Abstracts of papers in *Computational Intelligence*, 1(1), Spring 1985

Recovering from Execution Errors in SIPE

David E. Wilkins

In real-world domains (a mobile robot is used as a motivating example), things do not always proceed as planned. Therefore it is important to develop better execution-monitoring techniques and replanning capabilities. This paper describes the execution-monitoring and replanning capabilities of the SIPE planning system. (SIPE assumes that new information to the execution monitor is in the form of predicates, thus avoiding the difficult problem of how to generate these predicates from information provided by sensors.) The execution-monitoring module takes advantage of the rich structure of SIPE plans (including a description of the plan rationale), and is intimately connected with the planner, which can be called as a subroutine. The major advantages of embedding the replanner within the planning system itself are:

1. The replanning module can take advantage of the efficient frame reasoning mechanisms in SIPE to quickly discover problems and potential fixes.

2. The deductive capabilities of SIPE are used to provide a reasonable solution to the truth maintenance problem.
3. The planner can be called as a subroutine to solve problems after the replanning module has inserted new goals in the plan.

Another important contribution is the development of a general set of replanning actions that will form the basis for a language capable of specifying error-recovery operators, and a general replanning capability that has been implemented using these actions.

Plan Parsing for Intended Response Recognition in Discourse

Candace L. Sidner

In a discourse, the hearer must recognize the response intended by the speaker. To perform this recognition, the hearer must ascertain what plans the speaker is undertaking and how the utterances in the discourse further that plan. To do so, the hearer can parse the initial intentions (recoverable from the utterance) and recognize

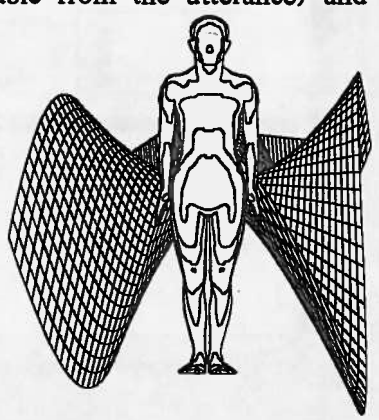

the plans the speaker has in mind and intends the hearer to know about. This paper reports on a theory of parsing the intentions in discourse. It also discusses the role of another aspect of discourse, discourse markers, that are valuable to intended response recognition.

Keywords: Discourse, task-oriented dialogues, intended meaning, speaker's plans, discourse understanding, plan parsing, discourse markers.

Knowledge Organization and its Role in Representation and Interpretation for Time-Varying Data: The ALVEN System

John K. Tsotsos

The so-called "first generation" expert systems were rule-based and offered a successful framework for building applications systems for certain kinds of tasks. Spatial, temporal and causal reasoning, knowledge abstractions, and structuring are among topics of research for "second generation" expert systems.

It is proposed that one of the keys for such research is *knowledge organization*. Knowledge organization determines control structure design, explanation and evaluation capabilities for the resultant knowledge base, and has strong influence on system performance. We are exploring a framework for expert system design that focuses on knowledge organization for a specific class of input data, namely, continuous, time-varying data (image sequences or other signal forms). Such data is rich in temporal relationships as well as temporal changes of spatial relations and is thus a very appropriate testbed for studies involving spatio-temporal reasoning. In particular, the representation facilitates and enforces the semantics of the organization of knowledge classes along the relationships of generalization / specification, decomposition / aggregation, temporal precedence, instantiation, and expectation-activated similarity.

A hypothesize-and-test control structure is driven by the class organizational principles, and includes several interacting dimensions of research (data-driven, model-driven, goal-driven temporal, and failure-driven search). The hypothesis ranking scheme is based on temporal cooperative computation with hypothesis "fields of influence" being defined by the hypotheses' organizational relationships. This control structure has proven to be robust enough to handle a variety of interpretation tasks for continuous temporal data.

A particular incarnation, the ALVEN system, for left ventricular performance assessment from X-ray image sequences, will be highlighted in this paper.

Keywords: Knowledge Representation, Expert Systems, Medical Consultation Systems, Time-Varying Interpretation, Knowledge-Based Vision.

On the Adequacy of Predicate Circumscription for Closed-World Reasoning

David W. Etherington, Robert E. Mercer, and Raymond Reiter

We focus on McCarthy's method of predicate circumscription in order to establish various results about its consistency, and about its ability to conjecture new information. A basic result is that predicate circumscription cannot account for the standard kinds of default reasoning. Another is that predicate circumscription yields no new information about the equality predicate. This has important consequences for the unique names and domain closure assumptions. □

The keys to Artificial Intelligence are at Xerox.

The Xerox 1108 Series of Artificial Intelligence Workstations provide an affordable, high-performance line of personal computers in an integrated, interactive environment that greatly enhances programmer productivity.

- High resolution graphics display
- Interactive user interface
- Ability to mix interpreted and compiled code
- Multiprocessing capacity
- Support of an extensible, interpreted language
- Display-oriented programming tools
- Local area networks and data communications through XEROX ETHERNET
- 8 Mbytes virtual memory

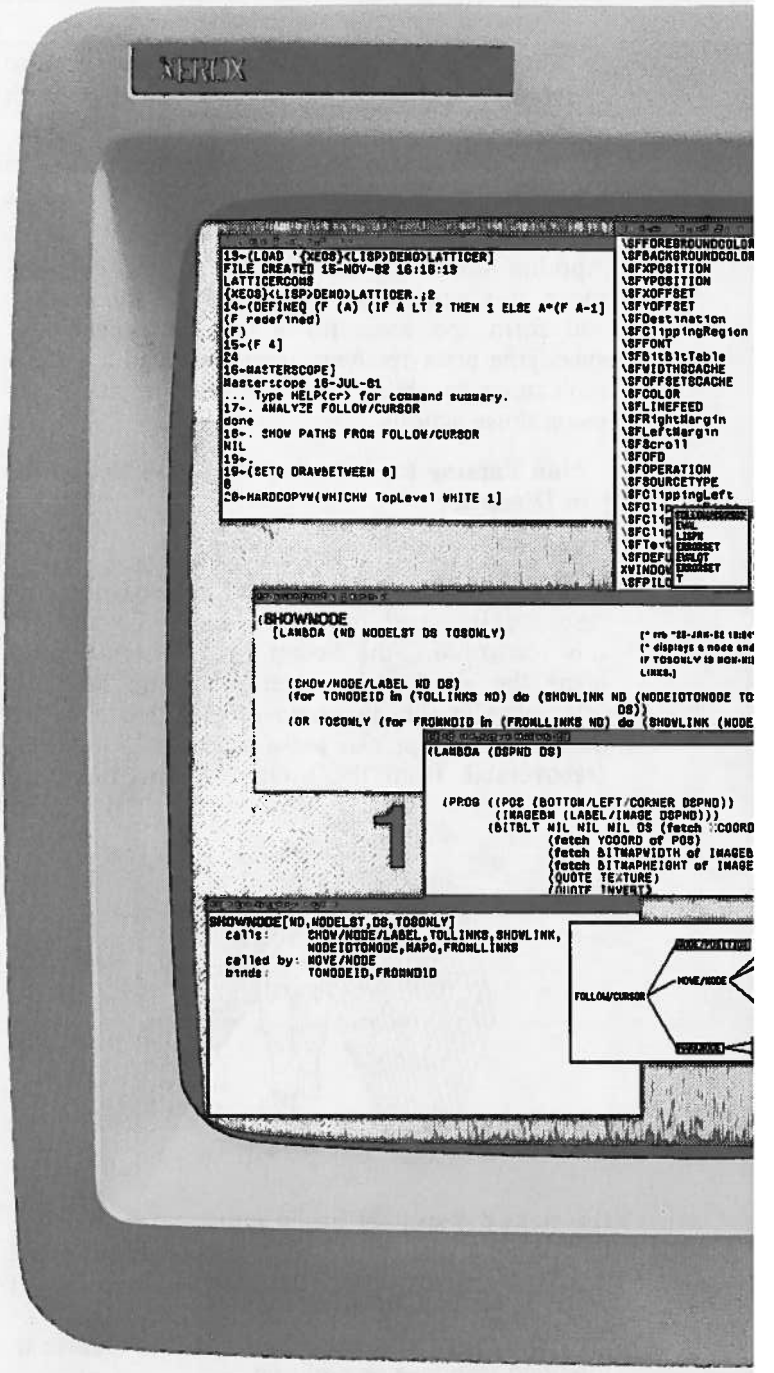

Power Tools for Programmers

1. Display Editor and Inspector

The display-based structure editor allows the interactive editing of programs and other list data. Structure-based editing exploits the form of an object, emphasizes the meaning of its parts, and thus reduces errors. The data inspector extends the philosophy to both system and user data types, allowing easy inspection and modification of any object in the system.

2. Programmer's Assistant

The Programmer's Assistant provides an intelligent assistant and bookkeeper that frees the programmer from much mundane detail. The Programmer's Assistant includes an error analysis capability and also monitors and records all user inputs. For example, a history is kept of the commands typed, their side-effects, and the results. Thus, one can request that a previous command or sequence of commands be repeated, modified and then repeated, or even undone (which undoes all the changes it may have caused). Also provided is a spelling corrector that automatically corrects spelling mistakes using information from the local context. To simplify file management for the programmer, Interlisp-D automatically keeps track of where in the file system each object is stored and which ones have been modified. In response to a simple request, the system can therefore save the user's state, updating all changed files automatically. The Programmer's Assistant provides a programming environment which cooperates in the development of programs allowing the user to concentrate on higher level design issues.

3. Debugging Tools

Debugging tools allow the user to break and trace

XEROX

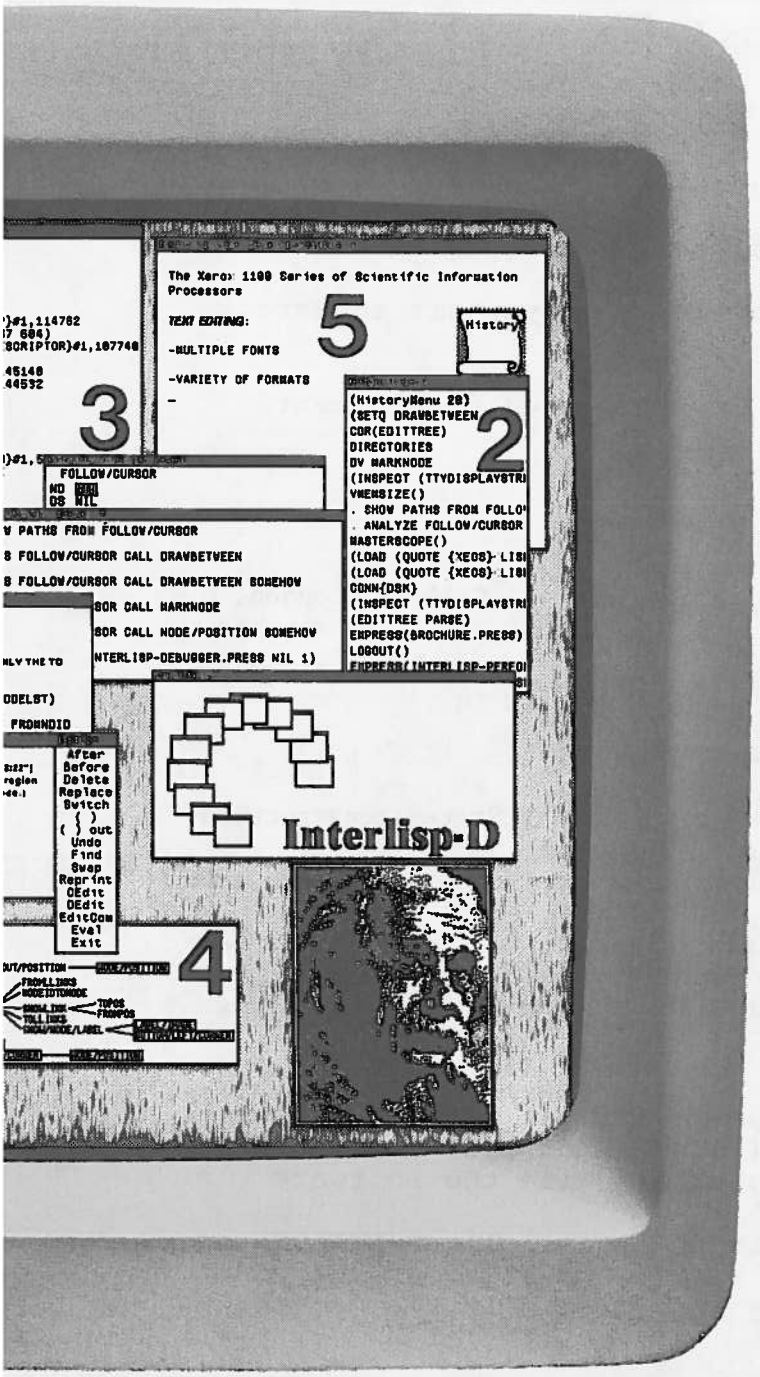

arbitrary functions, and examine the state of the machine at any desired level of detail. Not only can the state of a suspended computation be displayed and perused graphically, but it can be manually unwound to a specified point, the offending program edited, and execution resumed, all without loss of state. Also included is the capability of specifying complex, user-defined intervention conditions, such as allowing breaks only when a given function is called from another given function. These debugging tools allow bugs to be tracked down quickly and easily.

4. Program Analysis

The Masterscope facility can analyze a user's program and use that information to answer questions, display the program's structure and assist in the process of making modifications automatically. Because Masterscope is interfaced with the file package and editor, it re-analyzes a program whenever it is modified. Information about program calling structure, variable and data structure usage, and side effects can be graphically displayed and used to provide a map or browser for the system. The same information can be used to make systematic changes automatically. Further, Interlisp-D's measurement tools can be used to analyze the behavior of a system after it has been developed to pinpoint those areas that may need improvement.

5. A Professional Workstation

A high bandwidth user interface is provided by combining the mouse and the high resolution display. The mouse permits the user to specify and manipulate positions or regions on the screen. The interactive display facilities include complete raster graphic functions as well as a display management system supporting multiple overlapping windows, menu driven selection of operations, and a wide range of built-in graphical abstractions. Functions are also provided to display text in multiple fonts, manipulate raster images, and draw spline curves. The large format, high resolution display and the sophisticated multiple window system allow concurrent sessions, close-up views, and simultaneous displays of multiple representations of complex data. It is easy to create windows with text, graphics, or both and to make them scroll, update and interact in useful ways with the end user.

6. Knowledge Programming System (Optional)

LOOPS extends the programming environment to provide a powerful tool for research and expert system development. LOOPS combines four programming styles:

- Procedure-Oriented
- Data-Oriented
- Object-Oriented
- Rule-Oriented

Xerox Canada Inc.

Advanced Systems Group
703 Don Mills Road, North York, Ontario M3C 1S2 416-429-6750

Rapid Prototyping of Expert Systems

- Affordable introduction to in-house AI technology development
- Easy to use AI languages with several years' development
- An expert system "shell" to provide an easy start to Expert System development
- Optional technical support for Expert System development

- MICRO-PROLOG
  - Fifth Generation language for AI applications
  - Development began in 1975 (IC Prolog) at Imperial College, London, U.K.
  - Versions available for PC-DOS, MS-DOS, and CPM/86 operating systems
  - Basis of **apes** Expert System shell

- APES: AN AUGMENTED PROLOG FOR EXPERT SYSTEMS
  - Tools for effective logic programming and Expert System construction. Also developed by Imperial College
  - Runs on top of **micro-PROLOG**. Features of **micro-PROLOG** accessible from **apes**.
  - Simple Natural Language facility for rules, facts, and user interface
  - Tried and tested for 3 years: applied in Expert System construction in the domains of geology, law, biochemistry, medicine and engineering

- **micro-PROLOG** interpreter and **apes** available for $420.00 each. micro-PROLOG Reference Manual ($50.00) and apes documentation ($20.00) also available when purchased with the software
- An Expert System development kit including: **micro-PROLOG**, **apes**, micro-PROLOG Reference Manual, apes documentation, introductory book "micro-PROLOG: Programming in Logic", Clark & McCabe; available for $695.00
- All prices in Canadian dollars. PST not included
- Prices subject to change without notice (price changes expected in 1985)

Applied AI Systems Inc is an authorized dealer of **micro-PROLOG** and **apes**. micro-PROLOG and apes are registered trademarks of Logic Programming Associates Ltd and Logic Based Systems Ltd, respectively

AI
Applied AI Systems, Inc.
P.O. Box 13550
KANATA, Ontario
Canada K2K 1X6
Telephone: (613) 592-0084

Conference report

AAAI Workshop on Non-Monotonic Reasoning

David Etherington
Department of Computer Science
University of British Columbia
Vancouver, BC V6T 1W5

On 17 to 19 October 1984, a workshop on non-monotonic reasoning was held at Mohonk Mountain House, outside New Paltz, New York. The workshop was organized by Raymond Reiter and Bonnie Webber, and was sponsored by the American Association for Artificial Intelligence (AAAI).

The setting was spectacular. The hotel is an immense accretion of gables and towers in various architectural styles, facing onto a small lake. Surrounded by 2000 acres of private preserve, in full autumnal splendour, participants quickly forgot the outside world. The grounds include a small mountain, miles of walking paths, and 100 gazebos (*i.e.,* small parasol-bearing elk native to Catskill Mountain resort hotels).

The workshop program covered a variety of topics, with an emphasis on "formal" approaches to non-monotonicity (over half of the papers). As an exception to this trend, one participant, Gary Cottrell (University of Rochester), who presented a connectionist network implementation of inheritance hierarchies with exceptions, went so far as to describe himself as the workshop's "token scruffy".

Circumscription was a favourite topic, with 7 of the 22 presented papers dealing with different aspects and applications. John McCarthy unveiled a new, improved form of circumscription: "Formula Circumscription". Other papers discussed theoretical aspects of circumscription, as well as applications to default reasoning and discourse focus.

Default and auto-epistemic reasoning were also well represented, with a number of papers discussing aspects, applications, and implementations of default reasoning systems. Several papers emphasized non-monotonic facets of computational vision, natural language understanding, and common-sense reasoning.

An evening panel discussion was held, with John McCarthy, Dana Scott, and Richmond Thomason as panelists. The panel considered non-monotonicity from three perspectives: AI (McCarthy), Logic (Scott), and Philosophy (Thomason). Perhaps the most striking aspect of the discussion was the similarity of the three views presented.

Discussions continued at every available moment — during meals, hikes, rock-climbing, canoe trips, and in the hospitality suite — generally until late in the evenings.

The workshop's only disappointment was the shortness of time. Speakers (and the audience) often found that much more time could have been well-spent, especially considering the significance, quality, and complexity of much of the work presented.

A complete list of papers presented is given below. Preprints of the papers were distributed at the workshop, but no proceedings will be published. A limited number of copies of the preprints are available for ^{US}20 per copy from:

Claudia Mazzetti
AAAI
445 Burgess Drive
Menlo Park, CA 94301, U.S.A.

Papers presented

Nicholas Ascher (University of Texas, Austin)
Linguistic Understanding and Non-Monotonic Reasoning

Alexander Borgida and Tomasz Imielinski (Rutgers)
Decision Making in Committees — A Framework for Dealing With Inconsistency and Non-Monotonicity

Garrison Cottrell (Rochester)
Re: Inheritance Hierarchies with Exceptions

Jon Doyle (Carnegie-Mellon)
Circumscription and Implicit Definability

David Etherington, Robert Mercer, and Raymond Reiter (British Columbia)
On the Adequacy of Predicate Circumscription for Closed-World Reasoning

Martin Fischler and Oscar Firschein (SRI)
Computational Vision as a (Non-Monotonic) Reasoning Process

Clark Glymour and Richmond Thomason (Pittsburgh)
Default Reasoning and the Logic of Theory Perturbation

James Goodwin (Linkoeping)
WATSON: A Dependency Directed Inference System

(Continued on next page)

Benjamin Grosof (Stanford)
Default Reasoning as Circumscription

Joseph Halpern (IBM) and Yoram Moses (Stanford)
Towards a Theory of Knowledge and Ignorance: Preliminary Report

Aravind Joshi, Bonnie Webber (Pennsylvania) and Ralph Weischedel (Delaware)
Default Reasoning in Interaction

Vladimir Lifschitz (Stanford)
Some Results on Circumscription

Witold Lukaszewicz (Warsaw)
Considerations on Default Logic

W. Marek (Kentucky)
A Natural Semantics for Modal Logic Over Databases and Model-Theoretic Forcing I

Joao Martins (Instituto Superior Tecnico, Lisbon) and Stuart Shapiro (SUNY Buffalo)
A Model for Belief Revision

John McCarthy (Stanford)
Applications of Circumscription to Formalize Common-Sense Knowledge

Thorne McCarty (Rutgers)
Programming Directly in a Non-Monotonic Logic

Jack Minker and Donald Perlis (Maryland)
Protected Circumscription

Robert Moore (SRI)
Possible-World Semantics for Autoepistemic Logic

Mary-Angela Papalaskaris and Alan Bundy (Edinburgh)
Topics for Circumscription

Donald Perlis (Maryland)
Non-Monotonic and Real-time Reasoning

David Poole (Waterloo)
A Logical System for Default Reasoning

James Reggia, Dana Nau (Maryland)
An Abductive Non-Monotonic Logic

Donald Perlis (Maryland)
Bibliography of Literature on Non-Monotonic Reasoning

Renewal notices

Because CIPS handles administration for CSCSI/SCEIO, our renewal notices are printed on standard CIPS invoice forms. Don't let this confuse you — if you aren't also a CIPS member, you won't be slyly billed for a CIPS membership you didn't ask for.

Free Books

We have several AI books awaiting reviewers, who of course may keep the book after writing the review. If you would like to be a book reviewer for the *Canadian A. I. Newsletter*, contact the editor, giving your subfields of interest.

Social issues

Computer Professionals for Social Responsibility, Canada

Bill Havens
Department of Computer Science
University of British Columbia
Vancouver, BC V6T 1W5

Are you a computer professional who is concerned about society's use of computer technology? Computer Science, like Atomic Physics, is neutral with respect to its application either for or against humankind. Who should decide where computer technology is appropriate? Presently, computers are being used in all aspects of our society, and although the potential power and uses of computers can enhance our world, there are also inherent limitations and dangers in using them inappropriately.

The U.S. Defense Department intends to use AI and advanced computing technologies in a variety of weapons systems including strategic nuclear weapons. Will these systems increase or decrease our real security? As well, computers used in communications, databases, and control systems can present serious problems of privacy and safety.

We as computer professionals are the experts who best understand these issues. We cannot rely on the special interests of the defense establishment or private industry (in any country) to ensure our safety.

If you are interested in forming a Canadian society similar to the existing Computer Professionals for Social Responsibility (CPSR) in the United States, please write to me at the address above, or by electronic mail to havens@ubc (CSNET) or ubc-vision!havens (UUCP). □

New Bindings

Yawar Ali, from University of Toronto to Bell-Northern Research.

Timothy Bult, from University of British Columbia to Bell-Northern Research.

Book Reviews and Publishing News

Art and Computers:
The First Artificial Intelligence Coloring Book

Harold Cohen, Becky Cohen, and Penny Nii with a foreword by Edward A. Feigenbaum

William Kaufman, Inc., Los Altos, Calif., October 1984.
Distributed in Canada by John Wiley
(Hardcover, 128 pages; ISBN 0-86576-060-8; $^{CDN}28.95; $^{US}19.95)

Harold Cohen is an artist, who for the past sixteen years has been writing programs to produce drawings, which he subsequently colours. Naturally, that eternal question arises: "Yes, but is it art?" Before trying to answer, or even discuss this question we first turn to the book itself.

The book is a well-produced, large format, volume. On the dust jacket and the endpaper are examples of finished works by the artist. Of the 127 pages, the first 13 consist of the usual title pages and a foreword by the well-known computer scientist, Edward A. Feigenbaum of Stanford University and Teknowledge. On the final 35 single-sided perforated pages are black-and-white computer-generated drawings, ready to be removed, coloured, framed, and hung. Now the meat of the book, as it were, are 43 pages, four of which serve as introduction to what is to follow by Penny Nii, a computer scientist at Teknowledge, and Becky Cohen an artist and photographer. The remaining 39 consist of a dialogue between Harold Cohen and two thirteen-year-old girls, Karin and Sheri Bryant.

This dialogue is ostensibly about how computers can produce such a variety of pictures. But many other issues surface: How do programs make decisions? What does it mean for a program to know about shapes and space? Can programs know anything? How do artists make decisions? All this is very heady stuff and the two young girls are by turn suitably skeptical, enthusiastic, and eager to please. The text is liberally sprinkled with photographs of the girls at work, sketches to illustrate programming principles, and fragments of possible program-produced pictures. The book itself can be read and digested in about 20 minutes, leaving considerable time for colouring. Three six-year-olds of my acquaintance enjoyed themselves colouring the pictures with little concern about either computers or AI.

Although the book is titled *The First Artificial Intelligence Coloring Book*, I would claim that this represents more of a marketing strategy than a fair description of the contents. Edward Feigenbaum's endorsement notwithstanding, the message is not AI but rather that computers can be fun and can do interesting things, hardly an earth-shaking message. AI in the title is meant to sell the book by appealing to the current high profile of AI in the public consciousness. Furthermore, the intended audience must be children for whom colouring is a typical and enjoyable activity. Would the same appeal be present without the high-tech connection?

It seems to me that this book is more successful as a brief exploration of the nature of art. For example, in response to one of the girls' ques-

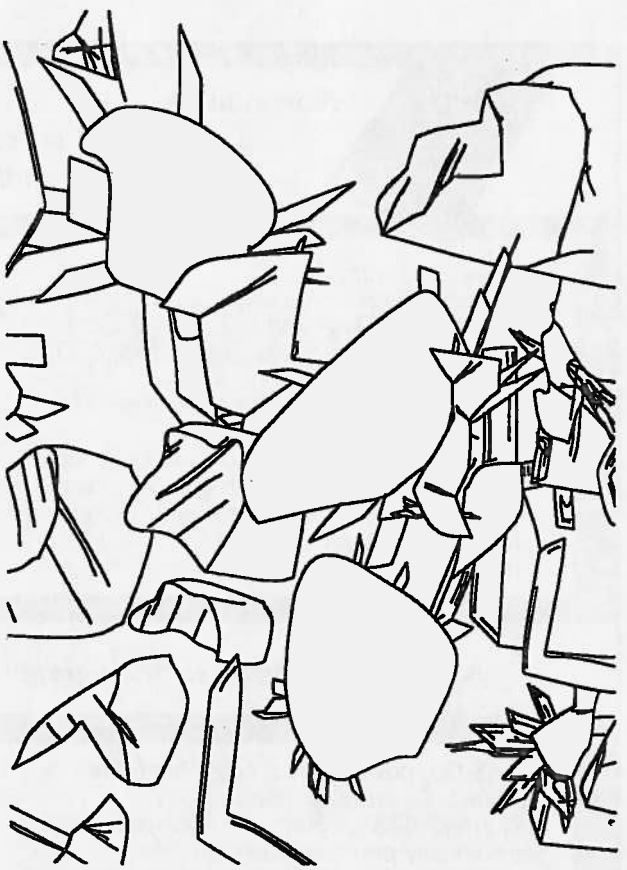

From *The First AI Coloring Book*

tions about how a simple drawing can invoke a particular feeling, Cohen responds:

> Well, in part it's because drawings look as if they were made by a person's hand, and you want to think that the person had something in mind to draw, so much so that you make something up. But it seems also that some kinds of marks — or some ways of making marks — let you make things up more easily than other kinds do. (p. 46)

Thus if the viewer plays an active role in the artistic process, which he or she does, then that role could certainly be assumed with an object produced by a computer. Our experience, our immersion in a cultural matrix equips and, indeed, encourages us to see, to understand, and to integrate visual experiences into our everyday activities.

This book, with its pretentious title and its padded format, nevertheless appeals on two levels. For children, it provides the opportunity to colour abstract diagrams and to create their own world. Of course, there are less expensive ways to amuse children. For the more mature, there is a brief, somewhat superficial, introduction to knowledge and decision-making in programs, and to the relation between computers and art. In an age in which exercise is organized in 20-minute segments, perhaps it is only natural that books are similarly packaged.

— Richard Rosenberg
Dalhousie University

New journal: *Expert Systems*

A new quarterly journal, *Expert Systems*, has commenced publication. The editors are Don Waterman, Ian Croall, and Mitsuru Ishizuka. Subscriptions, at $US79 for 1985, may be ordered from the publisher, Learned Information Inc., 143 Old Marlton Pike, Medford, NJ 08055, U.S.A.

Report about *Cognitive Science* was correct

In the last issue of the *Canadian A. I. Newsletter*, it was reported that the journal *Cognitive Science* was rescinding the $US8.50 foreign postage surcharge for subscribers in Canada. Shortly thereafter, however, the Cognitive Science Society erroneously sent out renewal invoices that included the surcharge. Canadian subscribers should ignore the surcharge, and pay at the U.S. subscribers' rate.

Recent AI Technical Reports

Editor's note: Recent Canadian AI technical reports are listed in this department. Abstracts will be included as space permits, with preference being given to theses.

University of Saskatchewan

Requests for any of the following publications should be addressed to:

Gord McCalla
Department of Computational Science
University of Saskatchewan
Saskatoon, Saskatchewan S7N 0W0

LEPUS: An Object-Oriented Extension of Turtle Geometry

David Goforth

TR 83-9, M.Sc. thesis

Early attempts to use computers as aids to learning tended to emulate classical teaching strategies. Research aimed at discovering and developing the unique capabilities of computers for education is still restricted to the laboratory. Practical classroom systems are still many years away. An appealing middle-term solution is the LOGO system, which makes powerful use of an interactive computing environment while avoiding the difficult issues addressed by other research. In the LOGO philosophy, a learner first becomes a programmer by mastering the simple but powerful LOGO language and then uses LOGO to explore other domains of knowledge.

The graphics metaphor of controlling a 'turtle' as it moves around the screen of a terminal facilitates both the learning of the language and the application to other material. The metaphor applies well to precise and static domains such as geometry but is inadequate for the kinds of concepts typical of the social sciences and humanities.

LEPUS has been developed to extend the application of the LOGO philosophy to these domains. The major features added are a time factor to allow dynamic concepts to be modelled, a facility for defining many 'objects', in place of just one turtle, to allow interaction to be represented, and finally, a capacity to represent imprecise, complex and incompletely defined concepts.

LEPUS is inspired by the object-oriented languages but departs from them in its handling of interaction. The user gains complete control over the communication process including the sender, receiver(s) and medium of access.

An extensive example of LEPUS in a social science domain is included. This example illustrates the structure of a LEPUS application and shows how the system can be used in a learning environment.

Automated Concept Discovery in Data Structure Applications

Mostafa M. Aref

TR 83-10, M.Sc. thesis

Intelligent behaviour often involves some sort of discovery process. Recently, there have been several systems whose task is to simulate the process of scientific discovery, and to synthesize new theories in their domains. This thesis explores the nature of the discovery process, in particular viewing discovery as a search through spaces of data, rules, and concepts. The thesis concentrates on the design of a specific system which operates in the domain of data structure applications. The system consists of a large collection of heuristic strategies written down as heuristic rules and implemented in a program to carry out discovery in data structure applications.

The system demonstrates its ability to build new structures (e.g., ordered lists, trees, and binary trees), and develop new operations (e.g., binary search and search trees). The thesis discusses the behaviour of the system with variations in some of its features (e.g., decision criteria), the potential capability of the system to synthesize its own heuristics, and the possibility of applying the discovery program to itself.

An Architecture for Plan-Based Computer-Assisted Instruction

Darwyn R. Peachey

TR 83-11, M.Sc. thesis

After nearly two decades of development in the area of computer-assisted instruction (CAI), it is still very difficult and expensive to produce CAI programs; moreover, the instructional effectiveness of the resulting programs is often disappointing. Although the use of artificial intelligence methods in CAI (AICAI) appears promising, many problems remain unsolved. One of the weaknesses of current AICAI systems is a lack of

"global" knowledge about the course being taught. In a large, complex course, this can lead to an inefficient and confusing path through the course material.

The thesis proposes an architecture for building better AICAI programs by applying and extending techniques which were developed for planning and controlling the actions of robots. Detailed examples show how programs built according to this architecture are able to plan global teaching strategies using local information. Since the student's behaviour can never be accurately predicted, the pre-planned teaching strategies may be foiled by sudden surprises and obstacles. In such cases, the planning component of the program is dynamically reinvoked to revise the unsuccessful stategy, often by recognizing student misconceptions and planning a means to correct them. This plan-based teaching strategy scheme makes use of global course knowledge in a flexible way that avoids the rigidity of earlier CAI systems.

Generating Summary Responses to Natural Language Database Queries

Jugal K. Kalita

TR 84-9, M.Sc. thesis

Development of natural language interfaces to databases has been an active area of research in the past decade. Most of the natural language database interfaces developed to date emphasize the interpretation of natural language queries; relatively little has been done on the formulation of appropriate responses.

For a natural language query system to be truly natural, it must not only accept input in the language of the questioner, but must also produce responses that approximate those of human speakers. Hence, interactions with a natural language database system should have the properties and constraints normally associated with human dialogue (i.e., obey a principle of conversational cooperation). Recent research includes the implementation of natural language query systems having some such conversational features.

An important convention of human dialogue is that no participant monopolizes the discourse. Therefore, in a conversational exchange, it would be considered inappropriate for a speaker to respond with a lengthy list of data; a shorter, non-enumerative response is generally more appealing. Such "summary" responses obey the principle of cooperative conversation by avoiding verbosity. Furthermore, extensional responses can occasionally mislead the user by generating false implications where summary responses would not.

In this thesis, we investigate the problem of generating such discourse-oriented concise responses. We present the details of design and subsequent implementation of a system which produces such summary responses. The system undertakes a heuristic search of the data which satisfies a user's query expecting to discover any underlying implicatures. During this process, it receives guidance from the information stored in the knowledge base provided by the system builder.

Queen's University

Requests for any of the following publications should be addressed to the authors at:

Department of Computing and Information Science
Queen's University
Kingston, Ont, CANADA K7L 3N6

A Framework For Robotic Perception

Roger A. Browse and Susan J. Lederman

TR 85-165

The development of intelligent, flexible robots that can adapt to their environment will require sophisticated perceptual systems. There are drawbacks and limitations to the direct application of computational vision systems as the mainstay of robotic perception. A more promising approach is to consider the possibility of multimodal perception involving an active visual component, along with touch sensors using proprioceptive information. There are several benefits to this broader approach to robotic perception because the perceptual information derived from the different knowledge sources is complementary. There are several interesting and addressable questions about the interaction of the various knowledge sources and about the repercussions of active perception on robotic planning systems. In addition, other areas of psychological research can now be viewed as being directly relevant to robotic issues.

Programming Languages For Artificial Intelligence

Janice Glasgow and Roger A. Browse

TR 84-162
[To appear in the International Journal of Mathematics and Computers]

This paper briefly describes the two most popular programming languages for artificial intelligence applications: Lisp and Prolog. The capabilities and limitations of each language are reviewed in

the context of establishing the main requirements placed on artificial intelligence languages.

The Nested Interactive Array Language, Nial, is introduced as a language that combines logic and functional programming capabilities. Through comparisons with Lisp and Prolog, it is shown that the Nial system meets the basic requirements for artificial intelligence programming.

Université de Montréal

Requests for any of the following publications should be addressed to:

Département d'informatique et recherche opérationnelle
Université de Montréal
Montréal, Québec, CANADA H3C 3J7

Athena: un programme joueur de Scrabble duplicate

S. Chapleau et Guy Lapalme

#443, Février 1983
[paru dans *TSI*, 2(4), 1983, 249–256]

Cet article présente le programme Athena qui utilise un algorithme simple pour jouer au Scrabble[MD] duplicate. Il utilise une analyse guidée par le nombre de lettres impliquées dans chaque coup possible déterminé après chaque jeu. Une fonction de choix heuristique permet de réduire considérablement le nombre d'essais de mots. Athena est très rapide et obtient un score de plus de 98% de l'optimal.

Athena: a program for playing championship level Scrabble

S. Chapleau and Guy Lapalme

#464, January 1983

This paper describes a very strong program for playing duplicate Scrabble[TM]. This program plays at a championship level and wins: in tests, it has beaten world champions. It does so by playing almost optimally very simple kinds of moves and by using an evaluation function to cut the number of trials to be made. Athena is very fast and achieves a score of more than 98% of the optimal play.

Logicon: an Integration of Prolog into Icon

Guy Lapalme and S. Chapleau

#516, December 1984

This paper describes the coupling of logic programming with Icon, which is a programming language aimed at string processing. We show that Icon and Prolog have many similarities, and that their integration is feasible and desirable as the weaknesses of one can be compensated by the strengths of the other. In our case, a Prolog interpreter was written as an Icon procedure which can be linked and called by an Icon program. This interpreter deals with all Icon data-types and can be called in the context of the goal directed evaluation of Icon. Examples of use are given.

Intégration du langage Prolog au langage Icon

S. Chapleau

#156, Novembre 1984

Ce rapport décrit les points communs de Prolog et d'Icon: backtracking et génération de plusieurs solutions pour une même requête (Prolog) ou une même expression (Icon). L'implantation de l'interprète Prolog est ensuite décrite comme un ensemble de procédures Icon. L'intégration préserve l'indépendance des langages et offre trois modes de programmation illustrés par des exemples: Icon pur, Prolog pur et une combinaison harmonieuse des deux.

Simula 67 par l'exemple

Guy Lapalme et F. Bronsard

#144, Juillet 1984

Ce document rassemble une cinquantaine d'exemples de programmes Simula 67 qui est l'archétype de la programmation par objet. Chacun d'eux illustre un point particulier du langage.

Prolog comme aide à l'analyse du français

D. Goupil, R. Lacouture, Guy Lapalme, R. Masson, et L. Simard

#149, Mai 1984

Ce document regroupe deux rapports et décrit une banque de programmes Prolog-II. Le premier rapport décrit une approche pour la compréhension du groupe nominal en français et montre une application, programmée en Prolog, dans un domaine très restreint. Le second étudie la possibilité d'automatiser la compréhension de textes français en se limitant au vocabulaire et à la grammaire définis dans le cadre du projet "Français Fondamental". Un autre prototype Prolog illustre ces principes.

Le logiciel graphique SIMSEE et le logiciel d'animation SIMSEA

L. Langlois

#150, Août 1984

SIMSEE est un logiciel graphique orienté vers l'objet. Il permet de créer et de manipuler des objets graphiques en 3-D. SIMSEE est écrit en

Simula 67 et est utilisé via des programmes Simula 67 préfixés par la classe SIMSEE. SIMSEA est un extension de SIMSEE destiné à la production d'animations.

Génération de phrases à partir de réseaux sémantiques

Michel Boyer

#152, Mai 1984

Ce rapport décrit un système permettant de générer plusieurs phrases françaises synonymes à partir d'un réseau sémantique donné qui en décrit le sens. Ce système a été implanté et est opérationnel sur un domaine limité en grande partie par la taille du dictionnaire. Son originalité est qu'il est le premier générateur de phrases s'appuyant sur la théorie Sens-Texte de Melcuk et Zolkovskii.

A Brief Overview of Expert Systems and Knowledge Bases

D. Tsiang and C. Frasson

#154, October 1984

This report examines the original context of expert systems, and then details the concept through several definitions. It distinguishes the differences between expert systems and other types of sophisticated computer programs and leads to a set of basic characteristics. It also gives an overview of the current approaches to represent knowledge in the architecture of an expert system.

Compilation de programmes Prolog dans un environnement d'objets

Jean-François Lamy

#159, Février 1985

Les possibilités de programmation orientée-objet et les coroutines de Simula permettent une description élégante d'une machine virtuelle pour Prolog. Nous montrons comment réaliser l'équivalent des techniques utilisées par des implémentations sophistiquées (traitement de la récursivité de queue, pré-traitement de l'unification, indexage des clauses), tout en demeurant dans le contexte de la programmation orientée-objet. La disponibilité d'un ramasse-miettes et la flexibilité des coroutines permet d'envisager des stratégies de fouille plus élaborées.

University of Toronto

Requests for any of the following publications should be addressed to:

Artificial Intelligence Group
Department of Computer Science
University of Toronto
Toronto, Ont., CANADA M5S 1A4

A Survey of Approaches for Determining Optic Flow, Environmental Layout, and Egomotion

John Barron

RBCV-TR-84-5
November 1984

Optic flow is the 2-D velocity field observed on an image plane that arises from the relative motion of 3-D objects and an observer. Five approaches to the computation of the optic flow field are discussed: template matching, difference techniques, token matching, gradient techniques, and velocity selective mechanisms.

We also survey approaches that attempt to solve the *reconstruction problem.* They use the knowledge of how *projective transformations* form a monocular 2-D image sequence from a 3-D scene. Given the optic flow field, these approaches attempt to invert the imaging process to compute *egomotion*, (the observer's motion), as well as the motions of the 3-D objects in the scene, and the *environmental layout* (the relative depth and orientation of each surface in the scene). Three approaches to the reconstruction problem are discussed: use of distinct points, local flow fields, and continuous flow fields.

All approaches are analyzed with respect to their feasibility, usefulness and applicability to 2-D imagery. Research is (usually) presented in chronological order with a subsection being devoted to each author or research group. An extensive bibliography is provided.

A Cascaded Filter Approach to the Construction of Velocity Selective Mechanisms

David J. Fleet and Allan D. Jepson

RBCV-TR-84-6
December 1984

In this paper we consider the extraction of motion information from spatio-temporal visual signals using cascades of linear filters. The mechanisms presented are simultaneously well-localized in space-time and specific for a narrow range of image velocities. We show that, in principle, arbitrary degrees of velocity specificity and localization in space-time can be attained, and is restricted, in practice, only by spatial and temporal sampling rates. The linearity allows a straightforward frequency analysis to be used in the design and evaluation of particular instances

of these filters. The cascades allow a very efficient implementation in terms of both storage and computation.

The first layer of processing is motivated, in part, by some recent models of retinal processing suggesting inseparable behaviour in space and time. This inseparability appears to be useful for machine vision and the extraction of motion information. The subsequent layers of processing and the hierarchical computational framework in general are consistent with qualitative models of cortical processing in biological visual systems.

McGill University

The table below lists recent AI technical reports from McGill University. They may be requested from:

Ms. Fran Lew
Computer Vision and Robotics Laboratory
Department of Electrical Engineering
McGill University
3480 University Street
Montreal, PQ, CANADA H3A 2A7

Recent McGill University AI Technical Reports

Title	Authors	Number
Spatial pattern in sections of human muscle	K. Paton and Steve Zucker	83-1
The relation between proximity and brightness similarity	Steve Zucker, K.A. Stevens, and P. Sander	83-2
Controlling a Puma 260 robot from a VAX computer	G. Carayannis	83-3
Low-level image segmentation: An expert system	A.M. Nazif and Martin Levine	83-4
Mapping numbers into symbols	A.M. Nazif and Martin Levine	83-5
An optimal set of image segmentation rules	Martin Levine and A.M. Nazif	83-6
Rule-based image segmentation: A dynamic control strategy approach	Martin Levine and A.M. Nazif	83-9
The fox and the forest: Toward a type I / type II constraint for early optical flow	Steve Zucker	83-11
Deblurring Gaussian blur	R. A. Hummel, B. Kimia and Steve Zucker	83-15R
Robot manipulator control by acceleration feedback: A single-joint paradigm	J. Studenny and Pierre Bélanger	83-16
Receptive fields and the reconstruction of visual information	Steve Zucker and R. Hummel	83-17
The local structure of image discontinuities in one dimension	Y. Leclerc and Steve Zucker	83-19
Robots and hybrids	Martin Levine	84-2
Kinematic equations and solutions for McGill Vision and Robotics Puma 260	J. Lloyd and J. Studenny	84-3
Everything you wanted to know about OPS5 but were afraid to ask	A. Dill and W. Hong	84-5
The use of vision feedback in printed circuit board assembly	A-R. Mansouri and A.S. Malowany	84-6
Trajectory planning in time-varying environments, 1: TPP = PPP + VPP	K. Kant and Steve Zucker	84-7
Progress report, October 1983 to September 1984	—	84-10
Visual inspection of solder joints by computer	P.A. Merrill and Martin Levine	85-1

VICE PRESIDENT OPPORTUNITY

ARTIFICIAL INTELLIGENCE

Unitek Technologies Corporation is a small growth oriented company with an outstanding future. This is no pie-in-the-sky operation but one with the five-year plan mapped out and the cash flow and finances in place. To help achieve their aggressive plans we need an exceptional individual to lead a team of software developers in artificial intelligence and expert systems.

Suitable candidates will have:

- a track record at a senior management level
- experience in a high-tech environment
- a bottom-line orientation with an understanding of marketing
- the capacity to build a team, keep it motivated, and get things done on time

If the results are there the position will lead to equity and advancement. An excellent salary and relocation expenses are offered. Location—Vancouver.

Please reply in confidence quoting File No. 09220 to:

Western Management Consultants
Suite 1706 - 1177 West Hastings Street
Vancouver, B.C.
V6E 2K3

VANCOUVER
EDMONTON
CALGARY
SASKATOON
TORONTO

Western Management Consultants
AG Lennox & Associates

Queen's University
at Kingston, Canada

Computing and Information Science
Tenure-Track Position in Artificial Intelligence

We are looking for new faculty members to assist in expanding the Department's research, graduate studies, and teaching activities. We particularly solicit applications for a position in artificial intelligence.

The Department has seventeen full-time faculty supported by nine technical and three office staff. Current research activity includes computer vision, logic programming, robotic perception, the analysis and implementation of parallel algorithms and parallel architectures, cryptographic and systems aspects of computer security, programming languages, analysis of algorithms, medical image processing, interactive information systems, data bases, software engineering, and audio signal processing. We have approximately 40 full-time graduate students. Resource support includes an IBM 3081 system running VM/CMS in the Computing Centre, plus Departmental facilities that include several VAX-11's, NS32016 systems, a reconfigurable multi-processor laboratory, a Symbolics 3600, and a variety of other machines and special equipment.

Kingston is located midway between Toronto and Montreal near the Thousand Islands, one of the prime recreational areas in Canada. Housing costs are moderate and stable.

Candidates of both sexes are equally encouraged to apply; Canadian citizens and landed immigrants will be given preference.

Applicants should send a curriculum vitae and the names of three references by April 5, 1985 to: Glenn H. MacEwen, Head, Computing and Information Science, Queen's University, Kingston, Ontario, Canada, K7L 3N6. Telephone: 613-547-2915.

University of Toronto
Department of Computer Science

ARTIFICIAL INTELLIGENCE RESEARCH SCIENTIST POSITIONS

Applications are invited for several research scientist positions in the Department's Artificial Intelligence Group.

The Group has recently been awarded a major grant to purchase Lisp machines and to fund these research positions. The Group presently consists of six faculty, two research scientists, and approximately 30 graduate students, with interests spanning all of A.I.

Applicants in any sub-area of A.I. will be considered. Applicants should have an excellent proven research ability; they should have (or expect to soon complete) a doctorate, or have very extensive industrial research experience.

Duties may include research, teaching at the graduate level, and responsibility for day-to-day running of the Group's research projects.

The University of Toronto's Computer Science Department is large and well-established, and is highly-regarded throughout the world for its research. The Artificial Intelligence Group was described by the granting agency as "second to none in Canada, and of world class".

Apply in writing, with curriculum vitae and the names of at least three referees, to

Allan Borodin, Chairman
Department of Computer Science
University of Toronto
Toronto, Ontario
CANADA M5S 1A4

For more information, telephone John Mylopoulos (416-978-5180), John Tsotsos (416-978-3619), or Graeme Hirst (416-978-8747).

FIRST!

International Conference and Exhibition on Computer Aided Technologies

COMPINT 85
Sept. 9-12, 1985
Montréal, Canada

With the accent on Robotics, Artificial Intelligence and Computer Applications in Industry, COMPINT 85 will bring together academic and industry experts from North America, Europe and Japan for the first time.

The location

As the world's second largest French speaking city, Montréal presents a myriad of attractions to the international traveller: fine hotels, world famous restaurants, an underground city covering miles of boutiques, cinemas and restaurants, and Vieux Montréal, to name but a few.

Program highlights:

1. Invited Speakers

Dr. Raj Reddy
Director of the Robotics Institute at Carnegie Mellon University

Prof. Jonathan Allen
Professor of Electrical Engineering and Computer Science, and Director of Research Laboratory for Electronics, Massachusetts Institute of Technology.

Prof. David Parnas
Landsdowne Professor, Department of Computer Science, University of Victoria.

Prof. Toshiyuki Sakai
Head of Information Science Department, Kyoto University
President, International Association for Pattern Recognition

2. Special Sessions

A. Research activities on new generation computers.

A unique opportunity to meet representatives of the world's major projects on new generation computer technologies. The projects and institutions represented:

Strategic Computing Project (DARPA) USA

National Project for New Generation Computer Technology—Japan

Institute for Computer Research University of Waterloo—Canada
Prof. Eric Manning

Intelligent System Architecture (ESPRIT) Europe
Prof. Jean-François Omnes

ALVEY Program for Advanced Information Technology—Britain
Dr. Brian Oakley

B. Robotics and Industrial Systems with

Dr. Jean-Claude Latombe
Institut Mathématique Appliqué de Grenoble

3. Tutorials

Topics include:

- Artificial Intelligence
- CAD/CAM
- Local area networks
- Software engineering
- Robotic manipulators
- Robotic vision

4. Exhibition

An international exhibition, including some 220 display booths, will be dedicated to various aspects of computer aided technologies, from industrial through managerial, educational and scientific applications.

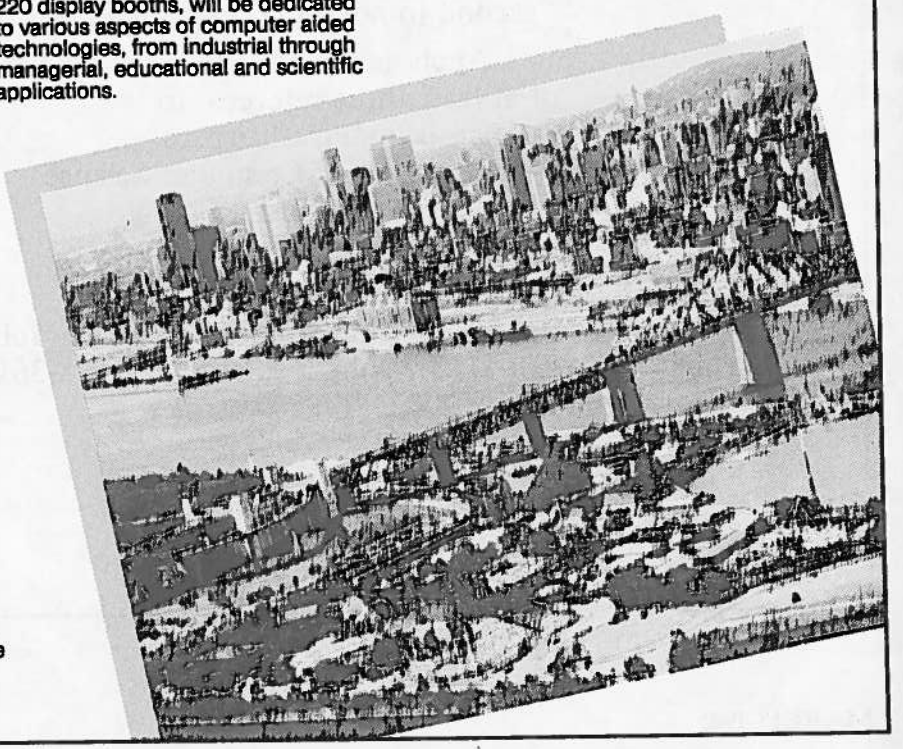

COMPINT 85
COMPUTER AIDED TECHNOLOGIES

For registration forms and more information, write to COMPINT 85, P.O. Box 577, Desjardins Postal Station, Montréal, Québec, Canada H5B 1B7 or phone (514) 394-5542.

acm

Forthcoming Conferences, and Calls for Papers

CSCSI/SCEIO Workshop on Theoretical Approaches to Natural Language Understanding

28–30 May 1985

Dalhousie University
Halifax, Nova Scotia

For details, see the announcement on pages 10–13.

Society for Philosophy and Psychology

15–17 May 1985

University of Toronto

This interdisciplinary meeting will include symposia and contributed papers on various topics of interest in AI, including computational theories of mind, mental states, memory and remembering. For more information, to register, and to arrange accommodation, contact:

David Olson
McLuhan Program in Culture and Technology
University of Toronto
Toronto, Ont., CANADA M5S 1A1
Phone: 416-978-7026

□ □ □ □ □

International Conference on Computers and the Humanities, 1985

26–28 June 1985

Brigham Young University, Provo, Utah

ICCH offers scholars and students an opportunity to meet and exchange ideas about their experience in using the computer as a tool for research and instruction in a variety of humanistic disciplines. Topics of AI-related interest at past ICCH meetings have included computer-assisted instruction, linguistics, language teaching, lexicography, graphics, textual analysis, translation, and word processing. For more information, write to:

Randall Jones / ICCH85
Humanities Research Center, 3060 JKHB
Brigham Young University
Provo, UT 84602, U.S.A.

□ □ □ □ □

15th International Symposium on Multiple-Valued Logic

28–30 May 1985

Queen's University
Kingston, Ontario

For more information, contact:

Prof H. T. Mouftah
Dept Electrical Engineering
Queen's University
Kingston, Ont., CANADA K7L 3N6

□ □ □ □ □

Workshop on Logic and Computer Science

9–14 June 1985

Lexington, Kentucky

This 4½-day workshop will cover those parts of Computer Science where an active part is played by logic-inclined researchers, in particular: theory of computation; theory of databases; artificial intelligence; theory of operating systems (temporal logic); program verification; logic programming. All inquiries should be sent to:

Logic and Computer Science Workshop
Department of Computer Science
University of Kentucky
Lexington, KY, 40506-0027, U.S.A.
Phone: 606-257-3961

□ □ □ □ □

Conference on Expert Systems in Government

23–25 October 1985

Tyson's Westpark Hotel
Tysons Corner, McLean, Virginia
(Suburban Washington, DC)

The conference objective is to allow the developers and implementers of expert systems in government agencies to exchange information and ideas first hand for the purpose of improving the quality of existing and future expert systems in the government sector.

Papers are solicited on expert systems in any domain. Original research, analysis, and approaches for defining expert systems issues and problems, methodological approaches for analyzing the scope and nature of expert system issues, and potential solutions are of particular interest. Send four copies of papers, no longer than 20 pages including graphics, by 1 May 1985 to:

Dr. Kamal Karna, Program Chairman

MITRE Corporation W852
1820 Dolley Madison Boulevard
McLean, VA 22102
Phone: 703-883-5866
ARPANET: Karna@Mitre

Notification of acceptance will be provided by 20 May 1985.

The conference is sponsored by the IEEE Computer Society and the MITRE Corporation in cooperation with The Association for Computing Machinery, The American Association for Artificial Intelligence, and The American Institute of Aeronautics and Astronautics National Capital Section. It will offer high quality technical exchange and published proceedings.

□ □ □ □ □

Conference on Theoretical and Methodological Issues in Machine Translation of Natural Languages

14– 16 August 1985

Colgate University
Hamilton, New York

The program of the conference will be biased toward invited lectures and panel discussions. However, a restricted number of excellent submitted papers will be also included.

The major topics of the conference are as follows :

- Machine Translation (MT) as an application area for Theoretical Linguistics (including stylistics and discourse analysis).
- MT as an application area for Artificial Intelligence (including the choice of the representation schemata for MT).
- Theory and methodology of translation and machine translation.
- Sublanguages, restricted domains and MT.
- MT as a case study in software system development.
- Computational tools for MT, human engineering aspects, management and evaluation of MT projects.

Papers should not exceed 3,000 words, and should contain a 250-word abstract and a list of index terms. The emphasis of the conference is on the theoretical and methodological issues. Therefore, the papers that do not address such issues will not be considered. Send papers (and address all inquiries) to

Sergei Nirenburg, Program Chair
Department of Computer Science
Colgate University
Hamilton, NY 13346, U.S.A.
Phone: 315-824-1000 x586

Submission deadline: 11 March 1985; notification of acceptance: 15 May 1985; final version due: 17 June 1985.

□ □ □ □ □

Cognitive Science Society 7th Annual Conference

15– 17 August 1985

University of California, Irvine

Topics will include: language processing, memory models, vision processing, belief systems, learning and memory, perception, knowledge representation, and inference mechanisms. For more information, contact:

Richard Granger
Computer Science Department
University of California
Irvine, CA 92717, U.S.A.

□ □ □ □ □

Conference on Space Station Automation and International Conference on Intelligent Robots and Computer Vision IV

15– 20 September 1985

Part of the Society of Photo-Optical Instrumentation Engineers' (SPIE) 1985 Cambridge Symposium on Optical and Electro-Optical Engineering

Hyatt Regency Cambridge
Cambridge, Massachusetts

In the next decade an increasing amount of research will be devoted to the applications of artificial intelligence and robotics technology to space station automation. The purpose of the conference is to bring together researchers in the areas of artificial intelligence, image science and robotics who are working on various aspects of space station automation. Papers on the following and related topics as applied to this unique microgravity, high vacuum, high radiation environment are invited.

Topics of interest include: space automation and tele-science; image understanding and scene analysis; machine vision; autonomous and self-organizing systems; hardware architecture designs; knowledge-based expert systems.

In addition, general papers on robotics and machine vision are solicited for the associated conference.

For information, contact:
Cambridge 1985
SPIE
P.O. Box 10
Bellingham, WA 98227-0010, U.S.A.
Phone: 206-676-3290
Abstract due date: 15 April 1985; manuscript due date: 19 August 1985.

□ □ □ □ □

International Conference on Genetic Algorithms and Their Applications

24–26 July 1985

Carnegie-Mellon University, Pittsburgh

Authors are invited to submit papers on all aspects of genetic algorithms, including the following topics: theoretical foundations of genetic algorithms; machine learning using genetic algorithms; classifier systems; apportionment of credit; genetic algorithms in function optimization and search; experimental applications.

Authors are requested to submit three copies (hard copy only) of a full paper by 1 May 1985 to the program chairman:
Dr. John J. Grefenstette
Computer Science Department
Vanderbilt University
Box 73 Station B
Nashville, TN 37235, U.S.A.

Morning sessions of the conference will be devoted to presentations of the accepted papers. Afternoon sessions will be devoted to panel discussions of the general themes raised in the morning sessions.

There will be no registration fee, but for planning purposes all attendees are asked to register by 1 June 1985. Registration information may be obtained from:
Dr. Stephen F. Smith
Robotics Institute
Carnegie-Mellon University
Pittsburgh, PA 15213, U.S.A.
Phone: 412-578-8811
ARPANET: sfs@cmu-ri-isl1

□ □ □ □ □

First Annual Workshop on Robotics and Expert Systems

27–28 June 1985

Johnson Space Center, Houston

The workshop is sponsored by the Instrument Society of America. Papers have been solicited in

the areas of expert systems, industrial robotics, computer algebra, AI for human–machine communications, teleoperations and space robotics, sensors and vision, distributed AI, and automated programming. Tutorials on expert systems and industrial automation will be offered. Registration is limited. For information, contact:

Dr Fred King
Ford Aerospace, M4B
PO Box 58487
Houston, TX 77258, U.S.A.
Phone: 713-280-6868

□ □ □ □ □

1985 Conference on Intelligent Systems and Machines

23–24 April 1985

Oakland University
Rochester, Michigan

Conference topics will include: intelligent robotics, machine intelligence, C^3I, adaptive control and estimation, visual perception and computer vision, pattern recognition and image processing, artificial intelligence for engineering design, intelligent simulation tools, computer-integrated manufacturing systems, knowledge representation, expert systems, game theory and military strategy, interpretation of multisensor information, automatic message understanding, natural language and automatic programming.

The conference will be preceded by tutorials on AI and Robotics held 22 April.

For more information:
Professor Nan K. Loh, Conference Chairman,
Center for Robotics and Advanced Automation
School for Engineering and Computer Science
Oakland University
Rochester, MI 48063, U.S.A
Phone: 313-377-2222

□ □ □ □ □

Fifth International Workshop on Expert Systems and their Applications

13–15 May 1985

Palais des Congrès, Avignon, France

For more information about the workshop, which will be conducted in both English and French with simultaneous translation, contact:

Jean-Claude Rault
Agence de l'Informatique
Tour Fiat – Cédex 16
92084 Paris – La Défense
France

Use this form to join CSCSI/SCEIO, and to order
Computational Intelligence and CSCSI/SCEIO conference proceedings

Canadian Society for Computational Studies of Intelligence

Société canadienne pour l'étude de l'intelligence par ordinateur

Application for Membership and / or Journal and Conference Proceedings Order

To join CSCSI/SCEIO and receive the *Canadian A. I. Newsletter*, fill out this form (or a photocopy of it) and send it to CIPS (which administers membership for the society) at the address below, with the appropriate fee. You need not be Canadian to be a member. This form can also be used to subscribe to the journal *Computational Intelligence* and to purchase CSCSI/SCEIO conference proceedings.

CIPS, 243 College Street (5th floor), Toronto, CANADA M5T 2Y1

Membership: $20 regular, $10 students (Canadian funds); there is a discount of $5 for CIPS members. *Computational Intelligence*: $16 / year (CSCSI/SCEIO members only). Conference proceedings: $25 each, plus $5 for postage within Canada, $7 for postage outside Canada. Payment may be made in U.S. dollars at the current rate of exchange.

( ) I wish to join CSCSI/SCEIO.

( ) I am a student. ( ) I am a member of CIPS.

( ) Please enter my subscription to *Computational Intelligence*.

( ) Please send me the following CSCSI/SCEIO conference proceedings:

( ) 1980 conference (Victoria). ( ) 1982 conference (Saskatoon). ( ) 1984 conference (London).

Total enclosed: $..... Cdn. / U.S.

Name .....

Mailing Address .....

.....

.....

.....

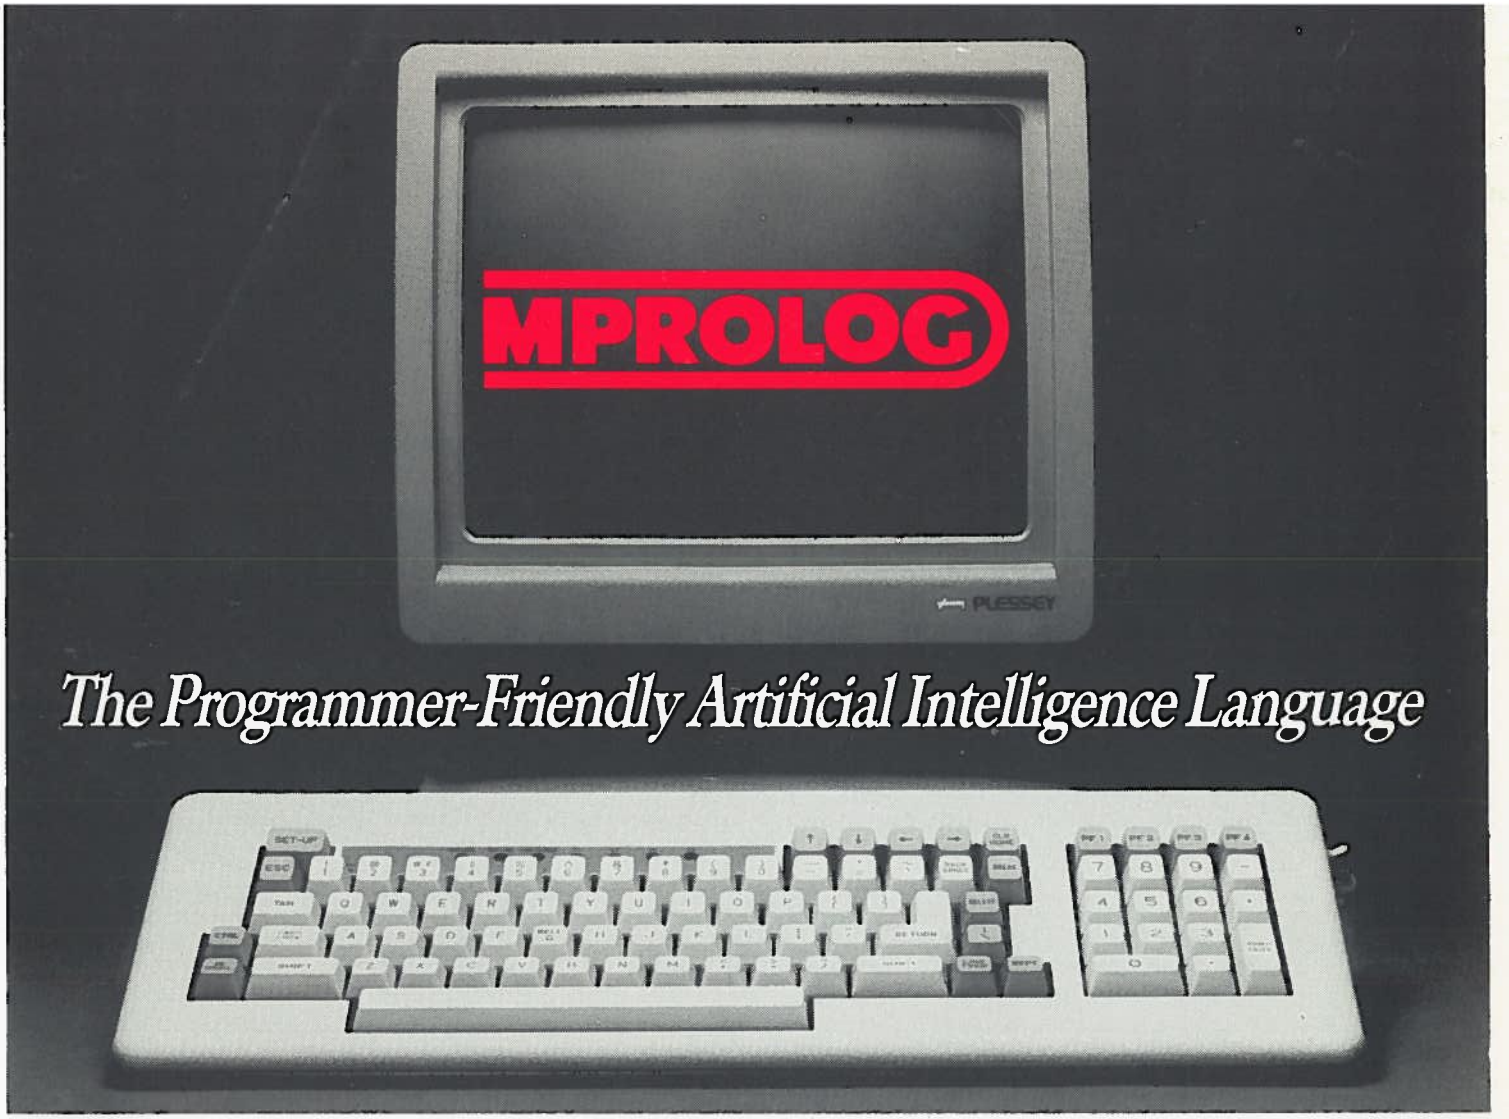

The Programmer-Friendly Artificial Intelligence Language

In what is unquestionably one of the most dramatic breakthroughs in the entire field of Artificial Intelligence (AI), a new modular version of the logic based programming language Prolog is now available exclusively through Logicware.

A refined, programmer-friendly version of Prolog — the computer language selected as the basis for the Japanese Fifth Generation Computer System Project — MPROLOG has a wide and powerful range of business, industrial and research applications.

With MPROLOG you can:

- protect the corporate investment in expertise
- create efficient expert systems,
- shorten programming development times
- provide transportability of applications

Why MPROLOG Was Developed

It has long been a dream of system developers to create a programming language which could be used to describe a problem without having to program explicit steps towards a solution.

A program whose execution corresponds to a **controlled deduction** through facts, relationships and rules.

A program that will enable computers to learn, associate, draw conclusions and make decisions.

MPROLOG is that program language, the language of artificial intelligence

When MPROLOG Can Be Used

MPROLOG can be used for:

- Expert Systems
- Deductive Databases
- Natural language understanding
- Computer-aided learning
- Fault diagnosis and repair
- Visual perception and guidance
- Intelligent Assistants
- Many areas of Artificial Intelligence

> "Day-to-day use of MPROLOG... by programmers working with specialists... over the past eight years has produced some 30 applications spanning everything from architecture to pharmaceutical research."
> The Financial Times, London (U.K.) November 13, 1983

MPROLOG Offers Seven Key Features

1. High performance with efficient use of resources
2. Hardware and operating system independence
3. Modular design capability; allows subsets of the problem to be specified and tested, offering substantial increases in productivity
4. Program development environment
   - Interactive Program Editor
   - On-line "help"
   - Concurrent editing and error correction
   - Program Trace
   - User-defined error handling
   - Efficient garbage collection
   - Over 250 built-in predicates
5. Interfaces to procedural languages and database managers
6. On-going support and enhancements
7. Comprehensive documentation and education

MPROLOG Works In A Flexible Operating Environment

The MPROLOG system is available on IBM mainframe computers under VM/CMS and MVS/TSO and Digital Equipment's VAX/VMS and VAX UNIX. Versions for the "super-micros" based on the M68000 CPU and the IBM PC-XT and AT are currently available with additional versions scheduled for release later this year.

To find out more about MPROLOG and Logicware call **1-617-547-2393**
(In Canada 1-416-665-0022)
or write

Logicware Inc., 1000 Finch Avenue W., Ste. 600, Downsview, Ontario Canada M3J 2V5

LOGICWARE

Boston • Toronto • Los Angeles

Towards a Canadian 5th Generation Research Plan

Canadian Society for Fifth Generation Research. An inventory of research expertise and a selection of proposed projects.

Drafted for the membership by Eric Manning, Nick Cercone, Ric Holt, John Mylopoulos, and Jean Vaucher, (Steering Committee).

Approved by ballot of the members, 21 January 1985.

March 1985. Supplement to the
CANADIAN ARTIFICIAL INTELLIGENCE NEWSLETTER

CONTENTS

TOWARDS A CANADIAN FIFTH GENERATION RESEARCH PLAN

1. Introduction

"The Fifth Generation will not be traditional computers. Instead, they'll be symbolic inference machines capable of reasoning their way swiftly through massive amounts of knowledge and data. They'll be computers that can learn, associate, make inferences, make decisions, and otherwise behave in ways usually considered the exclusive province of human reason . . . . Fifth Generation machines will understand spoken, written, and graphical input."

—P. McCorduck [*CACM*, **26**, 1983, 629–630]

The idea of Fifth Generation computers was sprung on the world in 1982 by the Japanese. Their Ministry of International Trade and Industry (MITI) announced an ambitious 10-year research program (with government funding of $500 million US dollars) aimed at developing a new generation of computing technology for the creation of machines that make logical inferences, deal effectively with masses of information, interact with humans using speech and images, and perform many tasks previously thought to require human experts. The launching of the Japanese research program was preceded by three years of study within MITI; the research project itself is staffed by a core of government scientists augmented by researchers seconded from Japanese industry. Owing to the low level of research activity in Japanese universities and to a squabble between MITI and Monbusho (the Ministry of Education) there is no real participation in the program by Japanese academics.

The Japanese program is heavily committed to the use of a computer language called PROLOG, which allows a programmer to state facts in terms of *logical assertions* such as

"John is the son of Tom",
"Tom is the son of Harry"

and *inference rules* like

"If A is the son of B and B is the son of C
then A is the grandson of C".

The running program is then able to answer questions such as

"Who is Harry's grandson?"

by making logical inferences. This kind of machinery is offered by the Japanese as a basis for building *expert systems* through the specification of assertions and inference rules used by human experts in a broad spectrum of applications ranging from the diagnosis of blood disease to the analysis of seismic logs. In the long run, the Japanese team seeks to design and build very powerful, general-purpose, *knowledge-based* systems which could make Japan the dominant force in the world computer industry. In the short run, they have already designed and built a small computer specifically tailored to execute PROLOG programs quickly and efficiently.

The potential of knowledge-based systems to transform almost every area of human activity — manufacturing, the resource industries, education, the professions, etc. — is reflected in the response of other industrialized nations to the Japanese challenge. In the UK, the *Alvey* program has provided funding of a similar level and scope to the Japanese, for research in microelectronics, software engineering, knowledge-based systems, and man-machine interfaces. The EEC, seriously concerned that its information technology industries could be given the coup de grace by Japan, has organized its own *ESPRIT* (European Strategic Programme on Research and Information Technology) program with committed funding of 11.5 million ECU and planned funding of up to 1.5 billion ECU over five years. In the United States we have seen the launching of three major programs. The *DARPA* (Defense Advanced Research Projects Agency) *Strategic Computing Program* seeks to maintain US world superiority in Fifth Generation Systems for both military and commercial purposes. DARPA plans to spend about $1 billion over the next six years to develop "intelligent" computing systems able, for example, to guide unmanned tanks and to assist military pilots. The money will be spent in US universities and industry to advance the technologies which

underly Fifth Generation computing: microelectronics, computer architecture, software engineering and Artificial Intelligence. Also in the USA, the *MCC* (Microelectronics and Computer Corporation) has been formed as a cooperative venture by several of the major computer and microelectronics firms to do advanced, long-term research in Fifth Generation technologies, while the *SRC* (Semiconductor Research Corporation) is a similar venture focusing on microelectronics.

In Canada, the Federal Government has not yet made a commitment to participate in the development of a Fifth Generation computer technology tailored to Canadian strengths and needs. Yet, without such a commitment, we will find ourselves increasingly uncompetitive in technology-intensive areas such as telecommunications and transportation where we now have a foothold to international markets. Also, our competitiveness in traditional areas of strength such as natural resources may be at risk, as the use of knowledge-based systems becomes increasingly widespread. Finally, Canada may be missing a unique opportunity to develop new software-related industries that would be able to compete internationally in already large and rapidly growing markets, if it doesn't take advantage of the strengths of research groups distributed across the country.

Hence this proposal.

2. Technologies and Applications

This section examines briefly the essential ingredients of Fifth Generation research — the *technologies* which underpin the drive towards Fifth Generation computers. In addition, the section mentions a few *applications* of Fifth Generation computing which are of particular importance to Canadians. Finally, it briefly surveys the *level of expertise* of Fifth Generation knowhow in Canada.

Fifth Generation computing has become feasible at this point in history because of the maturity of four key research areas: Artificial Intelligence (including topics such as natural language interfaces and knowledge engineering), computer systems and languages, microelectronics and computer architecture. The key task of creating Fifth Generation machines is the *integration* of achievements and techniques from these areas.

One of the major advantages of Fifth Generation systems lies in their *user interfaces* For example, *natural language front ends* will make it possible for untrained users to deal with computers in English or French. (Rudimentary natural language front ends are already commercially available.) Together with further advances in *speech recognition and synthesis* they will make it possible to converse with a machine as though it was a (human!) administrative assistant. The increased use of *computer graphics* will further improve the man-machine interface (the Apple Macintosh has already given the marketplace a glimpse of the advances which can be achieved here). Finally, the development of "user-friendly" human interfaces depends substantially on our understanding of the term "user-friendly". Research in *cognitive science* will give us a better understanding of the principles underlying "user-friendly" interfaces.

Turning to another type of interface between systems and their environments, *computer vision* has made striking advances in recent years. It will provide machines with the ability to "see" their environments and, more importantly, to "understand" what they are seeing. This capability will be essential for *robots* and will also allow machine interpretation of medical films to aid physicians with diagnoses.

Perhaps the most dramatic of the developments leading us to the Fifth Generation are in *knowledge engineering*, the set of languages, tools and techniques used to give a computer "knowledge" about a field of human expertise. The job of transferring knowledge, from the head of a human expert to the *knowledge base* of an expert system, has emerged as one of the most difficult part of the job of building expert systems. (The other parts of an expert system are its *inference engine*, which allows it to reason with the facts in its knowledge base, and the *explanation facility*, which allows it to explain its reasoning to the user.)

The knowledge base of an expert system may contain millions of facts. Storing them in a computer, retrieving them, protecting them from error or machine failure, sharing them among many users, and distributing them to many different machines — these are all difficult engineering problems. Progress in *computer systems and languages*, which includes work in software engineering, database management and programming languages, will allow us to address these problems adequately. The development of *programming environments* (a programming language plus all of the related software tools used by a programmer to design, specify, write, test and

document his work) is particularly important here.

Commercial viability depends heavily on the speed, cost and reliability of the product, and it is here that *microelectronics* and *hardware architecture* play a key role. Microelectronics, the technology of the "chip", allows us to shrink the physical size and cost of a piece of equipment, and dramatically improve its reliability at the same time. (The size and cost of electronic circuits has been dropping by a factor of ten every three years. By way of comparison, the difference in speed between a man walking and an automobile on a freeway is a single factor of ten, and the speed of a jet airliner represents another single factor of ten.) Researchers in hardware architecture are studying new ways to design and structure computers, so that many parts of a reasoning process can be carried on *simultaneously*. (Conventional machines are capable of doing only one thing at a time).

Turning to *applications* of Fifth Generation computing, there are several which will be crucial to Canada's future economic and social well-being. Advances in *robotics* are leading to mobile robots able to plan the execution of a complex task, see and understand their environment, and deal with unexpected obstacles. Examples could include robot miners, assembly line workers and lumberjacks. A second area is in *education*, where researchers have already developed systems to diagnose learning disabilities in children. Also, researchers are proposing a system for *Computer-Aided Learning* which builds an internal model of the concepts which the student does and does not understand; the system's interaction with the student is guided by this model. Yet another area is *Computer-Aided Design*, where Fifth Generation techniques will produce systems which will greatly enhance the capabilities and increase the productivity of human designers of electronic circuits, automobiles, metal parts, etc. We are pleased to note that this Plan contains proposals for substantial work in each of the application areas just mentioned.

Finally, what is the *level of expertise* in Canada? There is reasonable strength in all of the underlying research areas in our universities, mainly thanks to NSERC's policy of funding curiosity-driven research in a wide range of scientific and engineering disciplines. Indeed, some of the work required to develop Fifth Generation computing is under way already in Canadian universities. This is emphatically the case in Artificial Intelligence and Robotics, where the *Canadian Institute for Advanced Research* (CIAR) recently announced a national project based at three universities. It is interesting that the project is mostly privately funded and that it emphasizes high quality fundamental research. However, when we consider the breadth of the research areas which underpin Fifth Generation computing, the tremendous effort required to integrate these areas, and the size of other national efforts towards the same goals, we are forced to conclude that work is proceeding at too slow a pace and with insufficient cohesion.

The Canadian telecommunications industry and some of the domestically- controlled computer firms have expertise in varying degrees in some of these technologies. Also, firms in practically every sector of the economy are showing a lively interest in expert systems, natural language front ends and intelligent databases. (The fact that Canadian industry has provided unheard-of direct financial support to Canadian university research in Artificial Intelligence, through the Canadian Institute for Advanced Research, speaks eloquently of industry's interest in these topics.) Finally, there are a few fledgling firms whose purpose is to develop parts of Fifth Generation computing technology.

There is very little expertise in government in these areas; microelectronics work in the Federal Department of Communications, image processing expertise in Energy, Mines and Resources and a modest amount of expertise in robotics in the National Research Council provide the notable exceptions to this rule.

Appendix 1 of this report summarizes a recent survey of Artificial Intelligence research in Canada carried out by the Canadian Society for Computational Studies of Intelligence, and provides some documentation for the strength and breadth of Canadian research in one of the key areas of Fifth Generation computing.

3. Formation of the Canadian Society for Fifth Generation Research

Aware of the facts laid out above, about one hundred Canadian academics from universities from BC to Newfoundland met in Ottawa last March. About thirty participants from industry and government were also present. After a day-long meeting of position statements and discussions, the participants voted unanimously to form the

Canadian Society for Fifth Generation Research, in order to prepare a research plan for consideration by government. A steering committee was elected to prepare a draft; it was felt that the plan should:

- provide a balanced mix of short-term, quick-payoff projects and longer-term, more fundamental work;
- emphasize mission-oriented projects but provide *substantial* resources for the fundamental, curiosity-driven studies which lead to real, long-term advances;
- emphasize software and applications topics but certainly not to the exclusion of hardware projects which directly relate to other components of the Canadian Fifth Generation effort;
- emphasize those applications areas felt to be valuable to Canada, such as *natural resources, telecommunications, manufacturing, medicine,* and *education*;
- impose the discipline of the NSERC Strategic Grants Program, with its demands for precise statement of objectives and milestones, and the identification of a user for the research results. (This discipline is quite compatible with basic research as well as mission-oriented research).

Armed with these desiderata, the steering committee asked the community of academic researchers to prepare draft proposals for research projects. Because the *integration of knowledge* from several research areas is the key to Fifth Generation research, the committee encouraged *major proposals*, based on co-operation among several researchers or groups of researchers spanning several of the key research areas. In many cases this would require cooperation among two or more universities; the committee felt that modern electronic communication has made such cooperation feasible, and that it would be a most welcome addition to the academic research scene in Canada.

Twenty-seven proposals and submissions were received. (Some groups submitted descriptions of their expertise and activities in the field, but preferred not to come forward with definite budgets and objectives at present.) Several of them were large, multi-group proposals, and some involved multiple universities. They form the basis of the plan described in the next section.

4. Towards a Canadian Fifth Generation Research Plan

This section presents the proposals and submissions in two ways. First, it presents a classification of them according to the research areas which they emphasize, using the set of topics introduced in Section 2. We list and briefly describe each proposal under the heading that seems most relevant. (However, the reader should remember that many of the proposals address more than one topic.) Secondly, this section highlights a few "major initiatives" which are good examples of the type of ambitious, multidisciplinary proposal which the committee hoped to see. These examples are just that; they are not to be interpreted as proposals which the steering committee recommends for funding. (The steering committee was *not* given a mandate to adjudicate proposals; that task is the proper function of NSERC's peer review committees.) Short summaries of all relevant proposals are provided in Appendix 2.*

4.1. An Overview of the Proposals

Following the taxonomy of Section 2, we describe the proposals which emphasize Artificial Intelligence topics (namely interfaces and Knowledge Engineering), then Programming Languages & Systems, Architecture and Hardware (including microelectronics), and finally Applications. This grouping brings out some of the coherence which exists among the projects; there are also strong bonds across topics, some of which we mention in passing.

4.1.1. Interfaces

This area is key to making Fifth Generation systems truly convivial and easy to use. We received one proposal from a group at Simon Fraser University and the University of British Columbia, led by Veronica Dahl, to develop software tools based on PROLOG to assist in computer understanding, generation and translation of natural language. A closely related proposal was received from a group of seven researchers at l'Universite de Quebec, led by Lorne Bouchard, to develop software to support the analysis and production of written French. We also received a closely-argued plea from Richard Parent of

*Not included in this edition.

Universite Laval for high priority for natural language processing in this Plan. He noted the obvious applications to expert systems, Computer-Aided Instruction (CAI), translation and access to databases, and also argued that French-to-English translation is on one hand an important economic and cultural priority in Canada, and on the other hand a great opportunity, providing a possible market niche in supplying software to the francophone world market.

Finally, the emerging discipline of *cognitive science* has contributions to make to user interfaces and natural language understanding. Here we have a single but large proposal from Zenon Pylyshyn and several colleagues at the University of Western Ontario, to fund joint work by their Centre for Cognitive Science and Group for Computing Research. Projects proposed or underway include studies of the problems involved in transferring knowledge from human experts to computerized expert systems, vision research and studies of the human- computer interface. The basic method used is computer modelling of human cognitive processes. A proposal to develop an efficient, error-tolerant parser and logical interpreter for English is also a major component of a multifaceted plan received from Schubert and Pelletier at the University of Alberta.

Also related to the general area of natural language processing are proposals by Cercone and Colbourn. These are discussed below.

Turning to *computational vision*, there was one proposal, from Bill Havens of the UBC Laboratory for Computational Vision (LCV), which addressed this topic. He proposes to study "schemata" as a way to represent vision knowledge in computers, together with applications to *VLSI* (microchip) design and to the automated interpretation of aerial photographs. There was also a general submission from the LCV which included extensive documentation on the structure and activities of the laboratory, a major international centre for vision research. It is also useful to note the existence of another strong group in computational vision and graphics at McGill University, which is currently working on automatic inspection of printed-circuit boards via computer-vision technique, and whose work is included in the CRIM proposal. Finally, proposals by Wong on Robotics, Wortman on Knowledge Engineering and Pylyshyn on Cognitive Science include work which is closely related to computational vision. Moreover, DeMori's proposal includes research objectives related to speech understanding and character recognition, both important issues in the development of special purpose interfaces.

4.1.2. Knowledge Engineering

"Expert systems will be one of the key building blocks in the next generation of computer systems. The wide availability of expert systems will have a profound impact on the use of computers in industry, government and education. There is a strong possibility of a 'windfall' gain in productivity through the widespread introduction of expert systems in industry.

"It is certain that those countries which are actively pursuing Fifth Generation computing research will develop expert systems for industrial application. If Canadian industry is to have any hope of keeping abreast of this coming round of computerization of industry, it is imperative that we have the technology needed to create and use expert systems here in Canada."

—Adapted from Wortman et al.

A proposal by David Wortman and fifteen colleagues at the University of Toronto focuses on the development of a *framework* for building expert systems, based on research on knowledge representation, expert system control structures, linguistic and visual user interfaces, software engineering, databases, office systems and computer networks. The framework is intended to provide a complete environment for building expert systems, including powerful languages for knowledge representation, facilities for the construction of an inference engine tailored to a particular application, and a friendly interactive environment through which the knowledge engineer constructs his system. At the same time, the project proposes to address research issues on efficient knowledge retrieval, concurrent access, distributed knowledge bases etc.

Along different lines, Schubert and Pelletier at the University of Alberta propose to refine their notion of *semantic nets* as a superior way to represent knowledge, and to develop an associated language which offers better control over the temporal order of execution than is provided by the rigidly depth-first rule of PROLOG. This will be applied to support reasoning in a natural language understanding system and in a roving robot. We also received a proposal from Marlene Colbourn and two colleagues at the University of Waterloo to investigate a programming environment superior to PROLOG for the development of expert systems. Called CONLOG, it will attempt to model expertise as the *development of a theory*, updated as new facts are acquired,

unlike PROLOG which uses simple rules plus probability weights to model expertise. CONLOG will be evaluated by re-implementing an expert system, used to diagnosis learning disabilities in children, and comparing the re-implementation with an available implementation in LISP.

In addition, two groups proposed to build expert systems for specific purposes. Rudy Seviora and several colleagues at the University of Waterloo plan a system to perform some part of the *design of VLSI microcircuits*; and Subir Bandyopadhyay and John Devitt of the University of Windsor and University of Saskatchewan propose an expert system to *computerize scientific and engineering handbooks*. This system would produce appropriate formulas, equations and data in response to a user's requests, and would also provide symbolic computation by calling on existing systems such as Maple.

The proposals by Cercone and Van Emden, as well as the UBC submissions by Woodham and Havens of the Laboratory for Computational Vision, are also relevant here. Moreover, the Pylyshyn proposal includes knowledge acquisition as one of the research areas to be studied.

4.1.3. Programming Languages, Environments and Systems

The importance to Fifth Generation computing of new programming languages, such as LISP and PROLOG, was briefly sketched in the Introduction. Indeed, an essential factor in the explosion of interest in expert systems and knowledge engineering, has been advances in programming environments and the computer systems on which they run.

In this category, we received a proposal from Nick Cercone and six colleagues at Simon Fraser University and Queen's University, to develop and test a complete Fifth Generation system aimed at natural language understanding. To do so they plan to develop extensions to knowledge representation based on lambda calculus representations and lambda categorial languages, study *concept formation* in natural language understanding, design a new *language* based on NIAL, offered as a better implementation tool for AI applications than PROLOG or LISP, and finally implement the whole on both a multiprocessor and on a custom-designed cpu exploiting VLSI. The budget for this very ambitious project is typical of the larger projects proposed; it is about $1 miliion p.a. plus a single expenditure of $0.5 million for hardware. Of course, the figure quoted is not the true cost, in that it ignores the (at least equal) indirect costs which the Universities have been expected to cover.

Another research project with a similar goal is led by Tomasz Pietrzykowski at the Technical University of Nova Scotia. The goal of the project is to develop a new graphical logic programming system called LOGRAPH. The system subsumes PROLOG by offering several new language constructs and a way of viewing programs as pictures constructed on graphics terminals. Because of this, LOGRAPH is offered as a suitable programming tool for CAD and education applications.

Turning to proposals which focus on software engineering, Bob Probert, Gregor Bochmann and seven colleagues at the University of Ottawa and l'Universite de Montreal propose to develop an integrated set of Fifth Generation *software tools*, a workbench to help develop and check specifications, implement and test software, particularly software used in telecommunications. Bill Atwood and two colleagues at Concordia and McGill Universities propose to build a *distributed system* which can be used both to implement and simulate the distributed algorithms which are expected to figure prominently in future Fifth Generation systems. In addition, Wilf Lalonde and five colleagues at Carleton University plan to develop Actra, a machine based on active objects ("actors") that can communicate with each other in carrying out a task. The machine will utilize a multiprocessor architecture with separate processors for process scheduling, colour graphics and the like. The intended application area for the machine is industrial real-time Artificial Intelligence applications. As an alternative to the logic programming approach advocated by PROLOG, the actor approach emphasizes autonomous communicating experts which it realizes through an object-oriented message-passing paradigm.

Finally, a consortium of Montreal-area universities called the CRIM, (Centre de Recherche Informatique de Montreal, Inc) has submitted a description of CRIM, its composition, brief history and research goals, including a sketch of a proposal to create tools based on AI techniques for *CAD/CAM* (Computer-Aided Design / Computer-Aided Manufacturing) and for *office automation*. In CAD/CAM, work on computational vision, on expert systems to assist designers, and on numerical control of robots is proposed. In office automation, work on character recognition, data bases and expert systems is proposed.

Another proposal involving CAD was submitted by Mark Green and six colleagues at the University of Alberta, who propose to develop software tools to assist the construction of software packages for Computer-Aided Design; this is related to a proposal by Seviora to develop expert systems to aid the design of VLSI chips.

Finally, an umbrella project by Renato DeMori, Ching Yee Suen, and a number of colleagues at Concordia University, lists several research goals, including speech understanding and automatic learning, theory of problem solving and pattern recognition, distributed and parallel processing systems and algebraic and combinatorial computing. This research group is one of the principal components of the CRIM initiative (see section 4.2.1).

The proposals by McCalla, Wortman and Van Emden are also relevant to this section.

4.1.4. Machine Architecture and Hardware

We received a surprising number of proposals here, suggesting that the perception of Canadian researchers, as being excellent in software but uninterested in machine architecture and hardware, is perhaps in need of updating.

Forbes Burkowski and Randy Goebel of the University of Waterloo plan to build a workstation or *personal computer* tailored to the PROLOG language, thus potentially leading to a product which would compete directly with the first product from the Japanese project. Neil Ostlund of the University of Waterloo proposes a very large, powerful *PROLOG machine* (128 megabytes of primary memory, 2.5 gigabytes of disc storage) comprising 64 cells based on microprocessors. Finally, to complete this burst of hardware activity, Rudy Seviora and colleagues of the Waterloo VLSI group wish to define and build *VLSI chips* to provide architectural support for logic programming. The potential synergy with the projects just listed is clear. (The CRIM initiative and the Cercone proposal also include a hardware component, involving the use of VLSI technology in developing special purpose chips. All three are supported by the chip design and fabrication facilities offered to our universities by the *Canada Microelectronics Corporation*).

With more emphasis on software, Maarten van Emden and six colleagues of the University of Waterloo propose to build a suite of software, with the intent of developing a new personal computer for logic programming having the ease of use of the Apple Macintosh. Expert systems called *mundane expert systems*, to interpret well-documented bodies of knowledge such as government regulations, will be implemented on the new system.

In addition, the proposals by Cercone and Lalonde include tasks dealing with architectural and hardware design issues.

4.1.5. Applications

There were three submissions which focused on applications: Gordon McCalla and six colleagues at the universities of Saskatchewan, Waterloo, Simon Fraser and Western Ontario propose a knowledge-based programming environment which employs Artificial Intelligence, computer-assisted instruction and expert system techniques for the development of a *Computer-Aided Learning* system to teach programming. In particular, a computer model of the student's knowledge and of her misconceptions and un-mastered skills is built up; guided by this model a search of a set of rules is performed to determine what the system's next action should be. (This proposal is also noteworthy in that it demands a close working relationship among researchers in BC, Saskatchewan, and Ontario. This sort of interaction is now quite feasible, owing to the existence of *computer networks* which carry electronic messages and files.)

The second proposal was from Andrew Wong and colleagues of the PAMI Group at the University of Waterloo. They propose to construct a *mobile robot* (current industrial robots are stationary) capable of building up a picture of its environment as it roves around, and of using this picture to plan its actions (e.g. get to the lathe without crashing into the toolbox which someone just left on the floor). This project draws heavily on the basic work on knowledge representation, computational vision, and logic programming described throughout this proposal; the deliverable will be a functioning prototype of a roving robot.

Another submission by Bernard Moulin and three associates at Laval argues in favour of interdisciplinary collaboration and research on knowledge representation and knowledge engineering projects, in developing AI applications in social sciences, in the context of a French-speaking environment.

The proposal by Schubert also includes a plan for a roving robot capable of perceiving its

environment and planning collision-free sequences of actions.

4.2. Major Initiatives

Although the Fifth Generation projects in Japan, the USA and Europe have different thrusts, focus on different technologies and have set themselves different goals, they share some common characteristics. Prominent among these are an emphasis on both long term and short term research, and the encouragement of projects which span several of the disciplines which underpin the Fifth Generation. (Again, this is because the achievement of research goals for Fifth Generation computing is essentially a matter of *integration*; the harmonizing of results and achievements from several different topics in computer research.) Most important, Fifth Generation projects have encouraged research involving a *critical mass* of people, who often come from different organizations and even from different geographic locations. Noting this, the request for proposal which led to this report strongly encouraged initiatives that spanned several research topics underlying Fifth Generation computing. In this section we describe, in more detail, eight submissions that are consistent with this view, hoping that they will serve as examples of the type of research group or project which will be funded by the Canadian Program for Fifth Generation research.

4.2.1. The Centre de Recherche Informatique de Montreal, Inc.

The Centre de Recherche Informatique de Montreal (CRIM), which was created by the four Montreal Universities in October 1983, intends to become a leading Canadian centres for Fifth Generation research. CRIM's function is to actively promote and foster joint, concentrated research in computing areas vital to Canadian industry. At present, CRIM has the institutional backing of Concordia, McGill, Montreal, Polytechnique and UQAM, as well as the active support of over 50 researchers in the various constituent computer science and engineering departments.

Research in the topics which underlie the Fifth Generation is well established in Montreal, and some strong research groups are already in place. For example, McGill's computational vision and control laboratory was recently chosen as one of the sites for the national project on AI and Robotics of the Canadian Institute for Advanced Research. At the Universite de Montreal, the TAUM group pioneered automatic machine translation. Concordia is strong in the field of pattern recognition, particularly in the areas of voice and character recognition. At UQAM, researchers have developed a comprehensive package for the analysis of French text.

At present, CRIM has identified four primary application areas for research:

- *VLSI:* design tools and methodologies for reliable VLSI circuits including symbolic design, design verification, expert systems (Cerny et al.), design for self-testing (Agarwal et al.) and fault tolerant computing (Aboulhamid, Savaria et al.).
- *Intelligent Systems and Office Automation:* the design and application of knowledge bases, development of efficient inference engines, integration of logical, relational and object-oriented concepts and systems, optical character input; voice recognition, synthesis and storage; natural language processing (Bouchard, De Mori, Merrett, Vaucher, Kittredge et al.).
- *Telematique:* mixed voice and data, local area networks (Hayes, Hoang et al.); distributed systems and formal specification methods for automatic generation of protocols (Bochmann, Tropper, Kaplan et al.); videotext and mixed media databases (Gecsei, Houle, Johnson, Lustman et al.).
- *CAD/CAM:* development of AI tools for Computer Aided Design and Computer Aided Manufacturing based on expert design systems, computer vision and robotics (Zucker, Camarero, Patel, Thalmann, Levine et al.).

It is expected that 70-80 researchers will be involved in the projects, with approximately 15 working directly for the Centre. CRIM has already received grants totalling over $1 000 000,00 from both Federal and Provincial governments and a major portion of CRIM's revenue is expected to come from government sources but direct industry support is also being established.

The ultimate aim of the CRIM initiative is to combine resources of the Montreal universities and participating industry, through government financing, in order to establish a regional centre of excellence with strong capabilities in Fifth Generation research.

4.2.2. The University of Toronto Initiative

The University of Toronto has played a leading role in the development of computer science and computer engineering in Canada, through the establishment of several strong research groups within its departments of computer science and electrical engineering. The focal point of the proposed Fifth Generation project will be the Computer Systems Research Institute (CSRI). This organization, formed over fifteen years ago through an NRC negotiated development grant, has an international reputation in computer research. In the area of Artificial Intelligence, the work done at Toronto over the past decade has been recognized by the Canadian Institute for Advanced Research, which has made Toronto one of its three nodes for AI and Robotics research, with a focus on knowledge engineering. The goal of the Toronto project is to exploit synergistic relationships between its AI and other research areas, taking maximum advantage of existing expertise and previous research.

The key to the Toronto proposal is the interaction of researchers from several diverse areas. The rationale given for this approach is as follows:

- The AI group at Toronto has over a decade's experience in the design and implementation of *expert systems*. They are ready to move on to the next generation of expert systems, which will use more sophisticated knowledge representation and control structuring techniques, thus being able to handle new and more complex problem domains.
- A major component of every expert system is a large *knowledge base*. The design and implementation of the next generation of expert systems will require expertise in database-related technologies concerning efficient storage and retrieval of large amounts of data/knowledge, concurrency control, error recovery and distributed knowledge-based systems. This is an area in which the Toronto database group expects to make a unique contribution.
- Modern database and office systems are adopting some of the techniques used in the development of expert systems, such as knowledge representation techniques, planning and problem solving methodologies and natural language front ends for information systems.
- The next generation of expert systems will be very large computer programs. The techniques and experience that have been developed in *software engineering* will be necessary to manage the implementation and maintenance of these programs. Moreover, new compilation techniques will be required if the resulting software is to run efficiently. The software engineering group at Toronto plans to design special purpose software engineering tools to assist in the construction of large, efficient expert systems.
- In software engineering, the problem of *configuration management* (management of source and object modules in a large software system) is an issue of major concern, with significant economic importance. Correct and efficient configuration management is essential to any organization which produces large software systems. Another major software engineering concern is the design of *requirements modelling* languages to be used in the planning stages of a software project. One aspect of the Toronto proposal is the use of expert system and knowledge representation techniques to design a new generation of requirements modelling and software management tools.
- It is widely expected that the present trend toward *distributed computing* will continue and perhaps even intensify. It is clear that future expert systems will have to function within distributed computing environments. Another research team will contribute to the overall Toronto effort by studying *communication techniques and hardware* for future expert systems.

The consortium of professors involved in the proposed project is expected to involve more than 20 faculty members and 50 graduate students.

4.2.3. The Simon Fraser / Queen's University Initiative

Simon Fraser and Queen's Universities propose to integrate AI programming languages, AI systems software, machine architecture and VLSI studies to develop a Fifth Generation machine. This proposal is interesting both from the point of view of scope, dealing with issues ranging from natural language and computer vision research to machine architecture and VLSI design, and from the point of view of geography. The proposal brings together researchers based in British Columbia and Ontario to form the all-important critical mass for a project of broad scope. Queen's has a solid reputation in programming

languages and programming methodology and includes among its achievement the development of *NIAL*, an AI language which is in widespread academic use, and has already attracted commercial interest. Simon Fraser has a young and dynamic Computing Science department with a strong emphasis on AI and substantial experience in knowledge representation and natural language understanding, from both theoretical and practical viewpoints. SFU's experience in developing practical AI systems in their Laboratory for Computer and Communications Research (LCCR) will be used to evaluate NIAL implementations for AI systems, especially in comparison to LISP and PROLOG. The VLSI and hardware architecture research groups at both universities will address hardware issues concerning the construction of a Fifth Generation machine, including the design and construction of *custom VLSI* chips for a cpu tailored to AI programming. Both Queen's and Simon Fraser will use their expertise in robotic perception to ensure that the AI system to be developed can handle a range of AI problems extending at least from natural language understanding to robotic perception.

During the first three years, the project will concentrate on logical extensions to the NIAL programming language, intended to accommodate the knowledge representation and natural language understanding models developed at Simon Fraser University. The possibility of extending PROLOG, employing features from NIAL, will also be investigated. Moreover, architectures will be developed for implementing NIAL using new VLSI technologies and microprocessor environments. The architecture to be explored will attempt to effectively exploit the high level parallel description of problems possible in NIAL. Another aspect of the proposed project involves the implementation of NIAL on a parallel machine architecture, such as the BBN Butterfly, in order to achieve the speed of processing necessary for AI applications.

4.2.4. The Saskatchewan / Simon Fraser / Waterloo / University of Western Ontario Initiative

A project submitted by Gordon McCalla and six colleagues at the University of Saskatchewan, the University of Waterloo, Simon Fraser University, and the University of Western Ontario proposes to bring together concepts from Artificial Intelligence (specifically, AI approaches to Computer Assisted Instruction), Software Engineering, and Programming Environments. In particular, the project addresses the need to provide a comprehensive learning environment for any software engineering system, and an intelligent interface between the user and the underlying system.

The central thrust of the SCENT (Student Computing ENvironmenT) project is a fully-integrated knowledge-based programming / learning environment for Computer Science education. Phase I of the project deals with introductory programming. In addition to the usual capabilities of an interactive programming environment, such as editors, languages, debugging aids and the like, the SCENT environment provides an AI-based CAI subsystem that monitors a student's progress through a set of concepts to be learned as the student performs various tasks, and updates a student model to reflect the changing state of the student's knowledge. The student model is central to the SCENT approach, and governs all aspects of the interaction. In subsequent phases, SCENT will aim beyond introductory programming towards more general computer science and software engineering environments.

4.2.5. The University of Waterloo Institute for Computer Research

The Institute for Computer Research (ICR) at Waterloo is a federation of research groups and individuals pursuing computer research. It comprises eleven groups and a total of about fifty academics, together with about three hundred postgraduate students and a sizeable staff of professionals.

The ICR submitted seven inter-related proposals which collectively span a wide spectrum of hardware and software work. At the level of *hardware design*, Rudy Seviora proposes to define computational structures which would support high-performance logic programming machines, and to implement these as VLSI circuits. At the level of *architecture*, there are projects to design and build a prototype sequential PROLOG workstation, and a parallel PROLOG machine based on the systolic loop. Randy Goebel and Forbes Burkowski have outlined the design of a PROLOG workstation based on Sylvan, Burkowski's (multiprocessor) architecture which is intended to efficiently support message-passing software. The workstation will be based on a single Sylvan node; the current plan is to provide a micro-coded PROLOG kernel for an inference

engine and a conventional processor to execute database and I/O operations. Custom VLSI circuits can be simply attached to the message bus as they become available.

The facilities offered by the sequential PROLOG workstation will provide an environment for the further development of logic programming *languages and systems*, including Maarten van Emden's QUARFE proposal and Marlene Colbourn et al.'s CONLOG work. The QUARFE proposal envisions starting with conventional hardware, and providing an expert system-like user interface, for *all* interactions with a computer. The idea is to exploit the QUestion-Answer- Rule-Fact-Explanation conceptual interface of logic programming, to provide a sophisticated but conceptually simple user interface.

Colbourn, Poole, and Cohen plan to implement CONLOG, a logic programming language with extended logical and reasoning capabilities. CONLOG attempts to build a theory which is modified as additional asertions are acquired; this replaces the simple rules-and-probabilities structure of PROLOG. CONLOG will be evaluated by re-implementing an existing expert system and by developing a natural language interface based on Cohen's work on human discourse.

Andrew Wong and several colleagues in the ICR PAMI group will draw on their own expertise and the expertise of the other groups to construct a *roving robot*, capable of dynamically building up a knowledge base describing its environment (i.e. "learning" about its environment) as it moves around. This will be a major step forward in industrial robotics, particularly for flexible manufacturing cells. Wong's team have adopted the UBC group's work on the cycle of perception and will attempt to combine their existing techniques for fast knowledge-based pattern recognition with logic-based mechanisms for reasoning about a visual world. Finally, the VLSI group wish to study the application of expert systems to the problems of VLSI design.

Like the SFU/Queen's proposal, the ICR proposal is "vertically integrated", with work ranging from the silicon chip level to the natural language interface. The excitement felt by the researchers stems from this feeling of being able to tackle many of the major problems of a complete Fifth Generation system in an integrated, cohesive way.

4.2.6. The UBC Laboratory for Computational Vision

The objective of the Laboratory for Computational Vision is to achieve computer-based image understanding. The research strategy is to iterate the design cycle: propose, investigate and evaluate mechanisms for representing and using the knowledge necessary for vision to occur. The central concern is the design and implementation of effective algorithms that represent and exploit knowledge of the world and knowledge of how images are formed. This knowledge includes properties of sensors, laws of physical optics and information about possible configurations in the world.

At a practical level, the objective is to provide computer software which interprets images for particular domains such as satellite and aerial imagery of the earth's surface and sketch maps. This software is created as part of the design cycle in that the mechanisms proposed are embodied and tested in executable computer programs.

Computational vision requires the integration of many different forms of knowledge. A central research issue is how that large corpus of knowledge can be structured and used effectively. Furthermore, much work has been, and is being, done to make explicit the many implicit constraints for specific knowledge sources and to find good representations and algorithms for exploiting these constraints.

The purpose is to make machines see, partly to make them more useful and partly to learn more about the nature of perception itself. Computational vision is also providing a new language for algorithmic theories of human perception.

Faculty members in the laboratory are Havens, Mackworth and Woodham. Mackworth and Woodham are part of the UBC node of the CIAR program in AI and Robotics and collaborate with other members of the CIAR program in computer science and psychology. The laboratory also involves an interdisciplinary team of researchers affiliated with the UBC Remote Sensing Council (forestry, computer science, oceanography, soil science, geophysics and astronomy, geography, electrical engineering and civil engineering).

4.2.7. The University of Western Ontario Centre for Applied Cognitive Science and Technology

The University of Western Ontario has a well

established Centre for Cognitive Science, which works closely with members of the Departments of Computer Science, Psychology, Philosophy, Anthropology, Secretarial and Administrative Studies, the School of Library and Information Science and the School of Journalism. The proposal by Zenon Pylyshyn, director of the Centre of Cognitive Science, is for a new research centre which would combine the efforts of the Centre and the UWO Group for Computer research to study human cognition as it applies to problems of people using new generation computer technology. They also wish to study selected areas of AI, such as expert systems and intelligent sensors, which require basic knowledge of human cognition for their design. A fundamental aim of the Japanese Fifth Generation efforts is to make computers more accessible to all, and hence improve the quality of life. Such a goal, the proposal argues, cannot be achieved without a deeper understanding of human cognition and this should be included in the Canadian Fifth Generation program.

The three main research directions to be pursued by the centre are:

- Psychological and computational studies relating to knowledge acquisition. This involves research on expert behaviour in a range of tasks from medical diagnosis to decision making by executive assistants. The study will focus primarily on discovering a methodology for making the knowledge that experts use explicit, so that it can be transferred to the knowledge base of an expert system.
- Perceptual-motor research for sensory robotics and image encoding. This direction will attempt to bring together research on low-level vision and motor skills, currently carried out within the centre for cognitive science, to address issues that arise in extending the capabilities of industrial robots.
- Research and Development on the Human-Machine Interface. The research questions here can be grouped into two categories, depending on whether they involve experimental studies with human subjects using a computer system, or studies on how to design software interfaces which are easy to use.

One of the goals of the proposed centre is to produce more scientists who have a background in both computer science and psychology.

4.2.8. The Alberta Initiative in Design Environments and Intelligent Systems

Two teams of researchers at the University of Alberta have put forward proposals whose goals interlock in certain areas. A group of six researchers led by Mark Green has proposed to develop software tools to help the creators of CAD packages. The set of tools is sufficiently comprehensive that the term *design environment* is appropriate. Using this approach the team hopes to eliminate most of the programming currently required to produce CAD/CAM software. The use of AI techniques is planned and the results could be strategic for Canada, given the critical role of CAD/CAM in keeping our industry competitive.

The second team, comprised of Len Schubert and Jeff Pelletier, has proposed a multifaceted research program aimed at enhancing the inference capabilities of a logic-based semantic net developed over the past 10 years, greatly extending the natural language front end for this system, and augmenting the perceptual and planning capabilities of a roving robot. A general logical language of actions and plans is to be developed for use in the planning components of both the natural language system and the robot controller.

The research programs of the two teams will interlock in the areas of specification languages, expert reasoning, and natural language understanding. Both of the teams will benefit from a long-standing commitment to AI research at the University of Alberta, which has awarded many advanced degrees in this area and currently has 6 permanent faculty and 3 adjunct and visiting faculty with AI interests encompassing the proposals as well as computer vision, cognitive modelling, parallel search, and expert systems.

5. Observations

The total request for funds was for about $27 million operating plus about $6 million for equipment. Of the operating budget, $2.6 million was requested to buy partial relief (no more than half-time, and costed at about $5,000 per semester-course) from teaching loads, and $1.39 million was proposed to bring in senior research colleagues to participate in projects. All of the projects were costed for three-year duration, with one exception (van Emden, 6 years). The annual requested rate of about $11 million seems remarkably low in comparison with the Japanese figure

($50 million p.a. from government with at least as much again contributed by industry) or the Alvey, DARPA, and ESPRIT figures. (However, a substantial but as-yet-unknown additional sum should be added for communications and travel, to support the numerous multi- university proposals and to allow for effective maintenance of cohesion.) From another point of view, the requested amount represents less than a doubling of the amount now spent yearly by NSERC on university-based computer research in this country.

Current NSERC policy requires that the *indirect costs of research* be funded from other sources, and that NSERC funds not be used to buy teaching relief. Unfortunately, the universities are now simply unable to find other sources to cover indirect costs. Moreover, there is a serious shortage of computer science researchers, brought about by the tremendous international demand for individuals who can contribute towards future computer-related technologies, as well as the rapid growth of computer science programs at universities around the world. Typically, computer science professors are over-burdened with teaching, owing to very large numbers of undergraduate students and relatively small numbers of instructors. Hence a new approach must be taken if the research plan outlined in this report is to be realised. It is our view that the best course of action would be to provide *full funding* of indirect costs, and this is strongly recommended. Such a provision would raise the estimated annual cost of the proposed research program to approximately $20 million, and would be provided with the understanding that some of the overhead funding would be used to buy teaching relief for researchers. As a much less-desirable alternative, direct funding for *teaching relief* (for up to half the teaching load of a given researcher) could be considered, at a rate of $10,000 per semester-course. This would add $2.6 million to the three-year budget (teaching relief was priced at $5,000 per semester-course in the budget). However, it would provide Universities with $40,000 p.a. in return for four courses of relief. Four (semester) courses is the normal load for a faculty member and $40,000 is approximately the amount required to hire an additional junior faculty member.

Turning to the proposals themselves, there was a most encouraging trend to multi-university collaboration visible. There was also very good *balance* between quick-payoff and basic, long-term projects among the various areas of research which make up Fifth Generation computing and among the regions of Canada. The quality of the proposals was also generally high as far as can be judged from the preliminary nature of many of them. Moreover, most of the players whom we hoped to see did in fact submit preliminary proposals with very few notable exceptions.

6. Conclusions and Recommendations

Our conclusions are that:

- There is a plentiful supply of highly-qualified Canadian researchers in the fields which underly Fifth Generation computing. In fact, the quantity and calibre of people in Canada compares favourably, in our opinion, with the groups working elsewhere.
- The projects show an encouraging degree of coherence. Much of this coherence goes deeper and is more subtle than the examples we have been able to give in this report.
- appropriate, bearing in mind the success rate in competition, which may lie in the 30-50% range, and the additional requests for funds which do not appear in this report but will emerge once funding is available. (Incidentally, we wish to emphasize, one more time, that this document does not contain recommendations for the funding or otherwise of particular projects; such adjudication should be the province of NSERC committees.)

Our recommendations are as follows:

- Canada should have a program of research in Fifth Generation Computing; this document provides a basis for organizing the *academic component* of such a program.
- Serious consideration should be given to organizing a complementary *industrial component*. Such an organization is consistent with what other Fifth Generation projects have done and makes perfect sense in that it defines a bridge between research labs and production lines. The very substantial support provided by Canadian industry to academic research in Artificial Intelligence through the Canadian Institute for Advanced Research can serve as evidence of industry's interest in and support for this area of research. (Comments from Canadian industry, on the academic program proposed here, are provided in Appendix

3.*)

- The program should emphasize research in software and applications, but not to the exclusion of hardware research.
- The program should be administered through NSERC's peer review system.
- Funding for the program should account for direct as well as indirect costs; a less desirable alternative to a provision for indirect costs would be direct funding for teaching relief, at a rate of $10,000 per semester-course.
- Funding of the proposed Fifth Generation research program should be provided in accordance with the figures given in the previous section of this report; under no circumstances should this funding be provided at the expense of existing NSERC programs for computer science and computer engineering research.

Appendix 1: AI in Canada

Because Artificial Intelligence is a key topic underpinning Fifth Generation computing, a summary of Canadian expertise in AI is helpful. As we noted in section 2, most Canadian expertise in Fifth Generation computing lies in the universities, and AI is no exception. This account of Canadian activity in AI is based on a national survey carried out by Nick Cercone and Gordon McCalla on behalf of the Canadian Society for Computational Studies of Intelligence (CSCSI/SCEIO).

The survey found activity distributed from coast to coast, with approximately equal levels in the West and in the East. The major centres are Vancouver, Edmonton, Toronto, and Montreal. Most AI research is carried out in universities; however there is a small but growing number of firms starting up.

The major areas of expertise in AI, in descending order of popularity, are: expert systems and application, computational linguistics and natural language (NL) understanding, computer vision and image processing, and knowledge representation. There is also interest in AI programming environments, logic programming, cognitive science, problem solving, robotics, learning, educational applications, man-machine studies of AI, social impacts and theorem proving.

Two universities, British Columbia and Toronto, can be identified as *centres of excellence* in AI because of the number of faculty members active in AI research, the numbers of postgraduate degrees produced, and the time duration of their activity. UBC's work centres largely on computational vision, especially as applied to the interpretation of satellite and photographic images. They also have expertise in knowledge representation and natural language understanding. Toronto's expertise is in medical expert systems, computer vision, knowledge representation, natural language understanding and applications to database and software engineering.

Four other universities have a major interest in AI. Simon Fraser University has activity in NL understanding (especially practical NL systems and parsing), knowledge representation (KR), logic programming and computational vision. The University of Alberta has work in KR, NL understanding, image processing, adaptive systems and cognitive science. The University of Western Ontario is active in cognitive science, AI programming, logic programming and theorem proving. McGill University is involved with computational vision and robotics.

Other groups with major AI projects underway include the University of Victoria (expert systems), the University of Calgary (man-machine studies), the Alberta Research Council (expert systems), the University of Saskatchewan (problem solving and applications in medicine), the University of Manitoba (robotics), the University of Waterloo (logic programming), Queen's University (computational vision and cognitive science), Carleton University (AI programming), University of Ottawa (logic programming and KR), the NRC laboratories (image processing), Bell-Northern Research (man-machine studies), Concordia University (speech understanding), l'Universite de Montreal (NL translation), Acadia University (logic programming) and Defence Research Establishment Atlantic (expert systems).

Respondents to the survey were asked to comment on desirable future directions for AI in Canada. Opinions shared by many of the respondents were as follows:

- Extra funds should be provided for AI-related research and development in the amount of $1 million to $10 million p.a. over the next decade.
- These additional funds should be disbursed primarily by NSERC, and secondarily by other government agencies and private sponsors.
- The funding should be provided for AI in conjunction with other areas (notably Computer Science), not AI alone.

*Not included in this edition.

- Funds are badly needed for equipment, teaching release, support staff and additional researchers.

As for strategies to be pursued over the next decade, the opinions held were the following:

- AI R&D should be widely distributed across Canada. (Views held here differed sharply between the non-academic (who favoured concentration of R&D in a few centres of excellence) and the academics (who held the opposite view).
- AI investigations should be directed at a wide range of topics.
- Both theoretical and practical work should be undertaken.
- AI research should be done in both universities and in industry.
- AI research should be largely aimed at long-term goals (an opinion strongly held by academics and much less strongly held by non-academics).
- AI R&D should emphasize software rather than hardware.

The general thrust of the plan outlined in this report is consistent with most of the views recorded above, and were also shared by most of the participants in the founding meeting of the Canadian Society for Fifth Generation Research.

Appendix 2: Project Summaries

[Not included in this edition.]

Appendix 3: Comments from Industry

[Not included in this edition.]