comp

Canadian Artificial Intelligence

Intelligence Artificielle au Canada

July/juillet 1988 No. 16

An official publication of CSCSI, the Canadian Society for Computational Studies of Intelligence
Une publication officielle de la SCEIO, la Société canadienne pour l'étude de l'intelligence par ordinateur

Program your way to the top with Gold Hill.

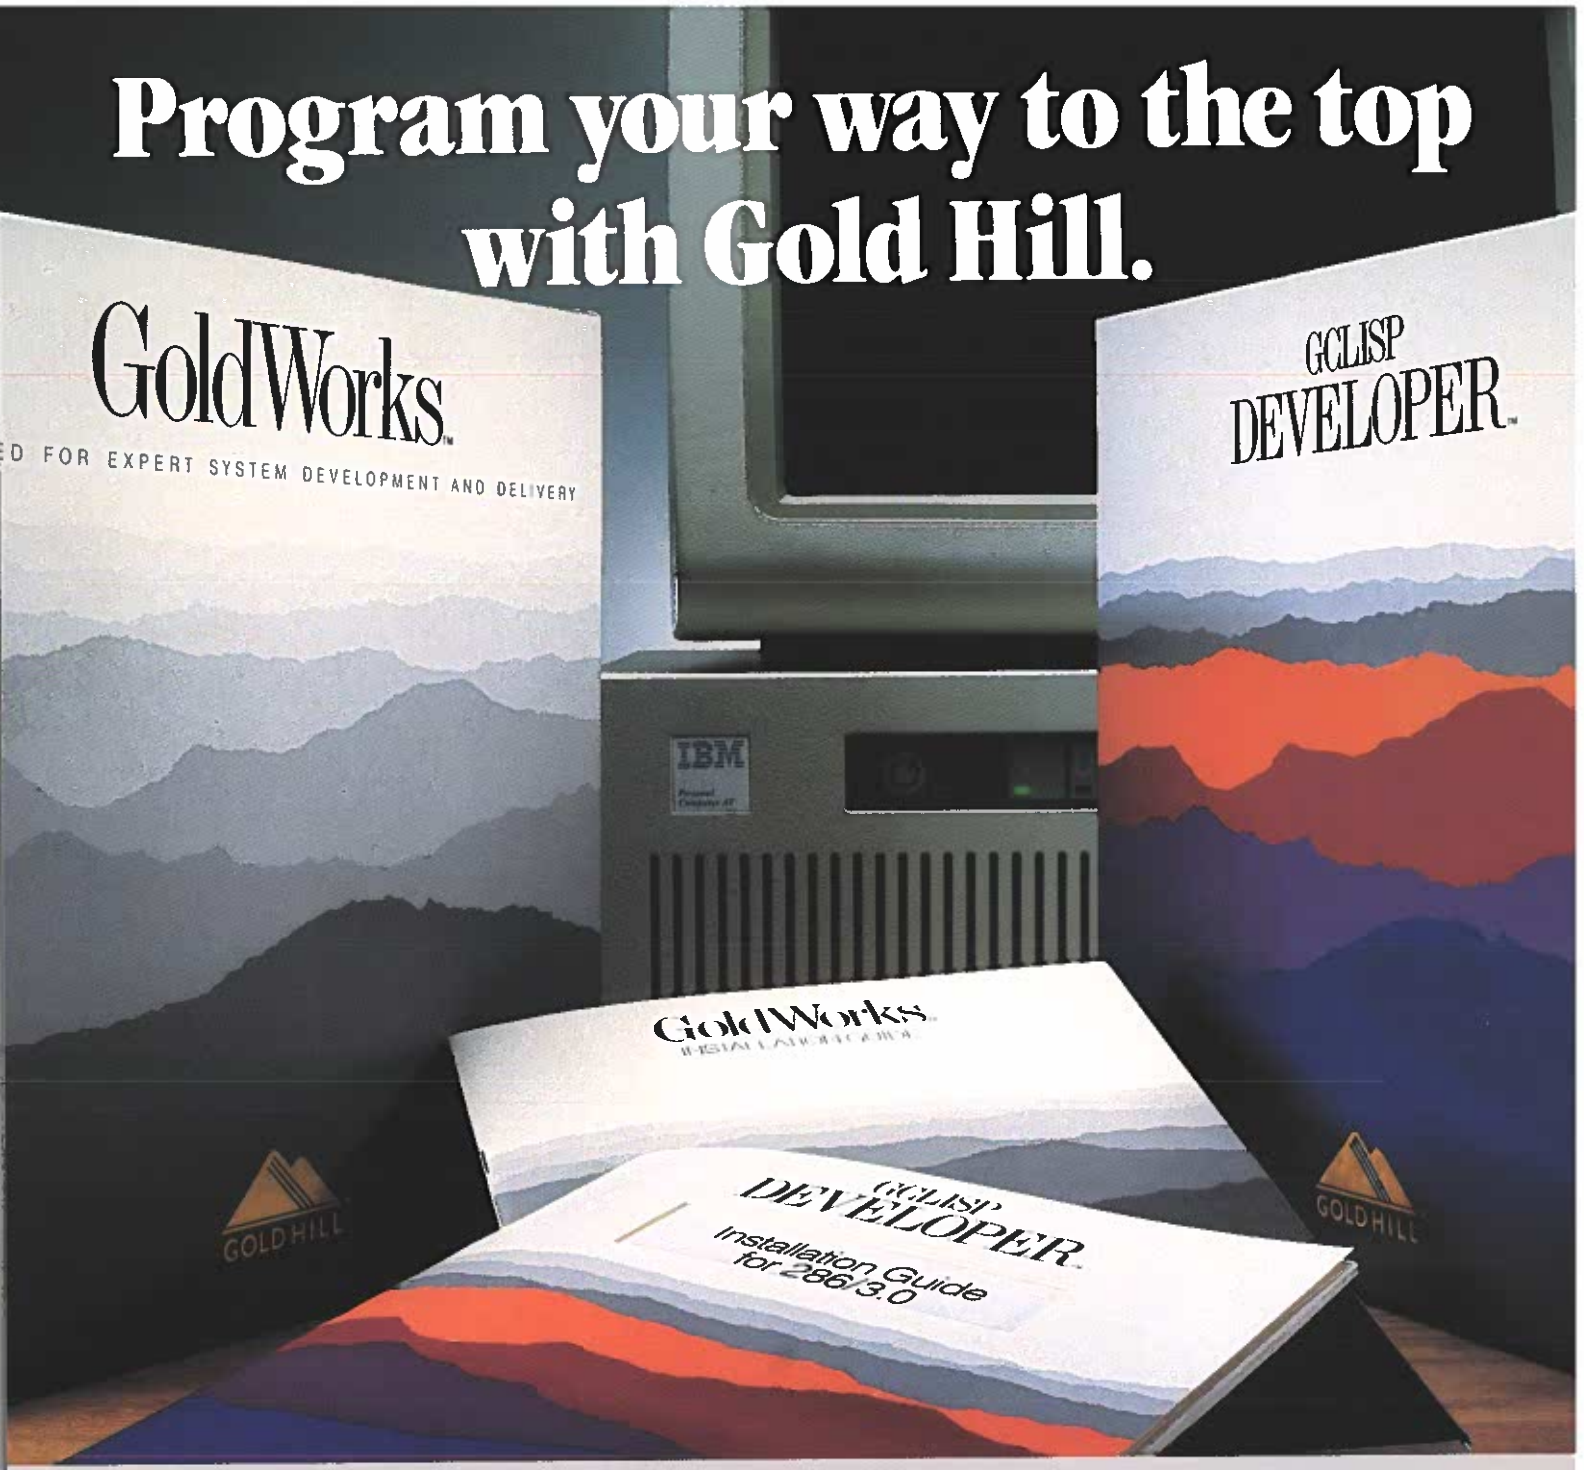

Your challenge is to develop the best expert system applications and deliver them to many users. Or to break new ground in AI development.

To succeed, you have to use the best development tools. They should be power tools that are technologically superior, cost-effective and proven. Tools that are fully supported with a comprehensive program of service, education and consulting.

Gold Hill's success depends on *your* success. That's why we're helping developers like you get to the top in AI—with **GoldWorks**, the PC-based expert system development/delivery tool, and **Golden Common LISP Developer 3.0**, the first *full* implementation of Common LISP on the PC under DOS. Products like these, plus an uncompromising commitment to excellence in service and support, make Gold Hill the number one provider of AI power tools for PCs.

If you're a programmer and you haven't taken a serious look at Gold Hill's power tools, write or call today.

Gold Hill. The expert in AI on PCs.

GOLD HILL™

Gold Hill Computers, Inc.
26 Landsdowne Street
Cambridge, MA 02139
Toll-free: 1-800-242-LISP
In Mass.: (617) 621-3300

Canadian Artificial Intelligence

Intelligence Artificielle au Canada

No. 16 July/juillet 1988

Contents

Canadian Society for Computational Studies of Intelligence

Founded 1973

CSCSI is the Canadian society for the promotion of interest and activity in artificial intelligence. It conducts workshops and fully refereed national conferences, publishes this magazine, sponsors the journal *Computational Intelligence*, and coordinates activities with related societies, government, and industry. To join CSCSI, use the membership form in this issue. Non-Canadian members are welcomed. CSCSI is affiliated with the Canadian Information Processing Society and International Joint Conferences on Artificial Intelligence, Inc.

Memberships in CSCSI:
Membership form is on the last page. Please send subscriptions, memberships, and changes of address to:

CSCSI, c/o CIPS, 243 College Street, 5th floor
Toronto, Ontario, CANADA M5T 2Y1

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

Fondée 1973

SCEIO est la Société canadienne encourageant l'intérêt et la recherche en Intelligence Artificielle. Elle organise des ateliers ainsi que des conférences nationales avec évaluation des articles soumis. Elle publie ce magazine, subventionne le journal *Intelligence Informatique*, et coordonne toute interaction avec des sociétés parallèles, le gouvernement, et l'industrie. Pour devenir membre de la SCEIO, veuillez utiliser le formulaire d'inscription de ce numéro. Les non-canadiens sont bienvenus. La SCEIO est affiliée à l'Association canadienne informatique, et aux International Joint Conferences on Artificial Intelligence, Inc.

Cotisations dans la SCEIO:
Le formulaire d'inscription est à la page dernière. Prière d'envoyer tout abonnement, cotisation, et changement d'addresse à:

SCEIO, c/o CIPS, 243 College Street, 5th floor
Toronto, Ontario, CANADA M5T 2Y1

Canadian Artificial Intelligence

Intelligence Artificielle au Canada

Founded in 1974 as / Fondée en 1974 en tant que *CSCSI/SCEIO Newsletter*

Submissions:
Canadian Artificial Intelligence is published quarterly by CSCSI/SCEIO and is a benefit of membership in the society. *Canadian AI* 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. Paper or electronic submissions are welcome. Electronic submissions are preferred and should be unformatted. *Canadian AI* is published in January, April, July, and October. Material for publication is due six weeks before the start of the month of publication.

Contributions:
L'*Intelligence artificielle au Canada* est publiée trimestriellement par la CSCSI/SCEIO, et est offerte gratuitement aux membres. L'*IA au Canada* encourage les contributions, en français ou en anglais, portant sur l'intelligence artificielle. Ceci comprend: des articles d'intérêt général; des descriptions de recherche courante et de cours; des rapports de conférences récentes et d'ateliers; l'annonce d'activités à venir, et des requêtes d'articles; des critiques de livres ainsi que des livres à critiquer; l'annonce de nouvelles compagnies en IA et de leurs produits; des opinions, des répliques, tout ce qui est polémique; des résumés de publication récentes, de thèses et de rapports; des trucs humoristiques ou artistiques, de bandes dessinées; des annonces (s'enquérir des frais); tout autre matériel touchant à l'IA. Contributions, sur papier ou par courrier électronique, sont bienvenues. Nous préférons le courrier électronique mais les submissions ne devraient pas avoir un format. L'*IA au Canada* apparait en janvier, en avril, en juillet, et en octobre. Toute communication à publier doit nous parvenir au moins six semaines avant le début du mois de parution.

Advertising:
Advertising rates and press kits are available upon request from the address below, or by phoning 403-297-2600.

Réclame:
Les prix pour les annonces et les trousses pour la presse sont disponibles sur demande. Écrivez à Marlene Jones à l'addresse à la gauche ou téléphonez 403-297-2600.

Please send submissions to / Prière d'envoyer contributions à:
CDNnet: CSCSI@noah.arc.cdn
UUCP: cscsi%noah.arc.cdn@alberta.uucp
ou à / or to: Marlene Jones
Canadian Artificial Intelligence
Alberta Research Council
6815 8th Street NE, 3rd floor
Calgary, Alberta, CANADA T2E 7H7

Book reviews and candidate books for review should be sent to:
Envoyez des critiques de livres ainsi que des livres à critiquer à:
CDNnet: gh@ai.toronto.cdn
CSNET: gh@ai.toronto.edu
UUCP: gh@utai.uucp
ou à / or to: Graeme Hirst, *Canadian Artificial Intelligence*
Department of Computer Science, University of Toronto
Toronto, Ontario, CANADA M5S 1A4

Canadian Artificial Intelligence is published quarterly by the Canadian Society for Computational Studies of Intelligence (CSCSI).
Intelligence Artificielle au Canada est publiée trimestriellement par la Société canadienne pour l'étude de l'intelligence par ordinateur (SCEIO).
Second Class Mail Registration No. 7373

ISSN 0823-9339

Executive Notes

Industrial Strength AI

It's in all the trade magazines — the big downturn in the AI business. A New York Times front page headline reads "Setbacks for Artificial Intelligence". The major corporate players are teetering on bankruptcy, being reorganized already or at least undergoing massive layoffs. Hardware sales of Lisp machines have saturated the market. Expert systems houses are continually losing money. What's going on? What happened to the glorious future of our emergent technology?

The critics have said it often enough before. Artificial intelligence is hard, very hard to achieve. The problems are more difficult than we thought; the representation issue unresolved; the lingua franca for the discipline undecided; the existing expert systems technology simplistic and brittle. Much theoretical research needs to be done before an industrial technology can possibly emerge from the laboratory. Our discipline has been oversold to industry by a few commercial charlatans. But what can be done about this malaise?

Push the genie back into the academic bottle. Let's retreat to our laboratories and work on the really hard problems some more. Eventually we will have the correct representational theories and from it a workable technology to apply to real problems. Until then, we should maintain a respectable distance from application issues.

No, we miss the point entirely. The emergence of artificial intelligence into the industrial environment is a natural evolution. More than that, it is necessary for the continued development of the discipline. The toy problems have been solved while the really difficult issues remain. How can we break the new conceptual barriers? What direction should our research take? What methodologies will yield new results? The answer is to work on real problems from which real insights can be found. Reality will force us to ask the right questions; to explore the right paths.

Recently, I had the opportunity to participate with our president, Dick Peacocke, in two seminars at Simon Fraser entitled "Industrial Applications of Artificial Intelligence". The audience was predominantly industrial computer scientists and managers. Both seminars were fully subscribed and the participants brought with them a receptiveness to new ideas and an enthusiasm for applying AI within their industries. Some already had fledgling expert systems projects underway. The seminars were judged to be very successful. Dick used the opportunity to encourage those present to join CSCSI and get involved in the Canadian AI community. I write this editorial to encourage you also to get involved, to get involved in industrial strength artificial intelligence. The opportunities are there. The problems real. And significantly, Canadian government and industry are willing to support collaborative research and technology development.

This is a transition point for our discipline and perhaps for CSCSI as well. Let's expand the scope of our organization by increasing membership from industry and government researchers. Let's use *Canadian Artificial Intelligence*, our biennial conferences and occasional workshops more effectively to engender communication and cooperation among academia, industry and government. Talk up our society with your neighbourhood industrialist. Take him/her to a power lunch and sign up a new CSCSI member. Find out which industrial applications can dovetail with your own research. The problems can be very interesting. Brainstorm a solution on the back of a serviette. Follow it up with a detailed research proposal. Look for joint industry/government funding. Get your graduate students involved. Get into Industrial Strength AI. It's an opportunity!

Bill Havens
Secretary

Humour—Dear Dr. Rob

Dear Dr. Rob:
Are there any questions you can't answer?

—Perplexed

Dear Perplexed:
Well, um, I'll have to get back to you on that one.

Dear Dr. Rob:
I have enclosed a copy of my recently developed AI tool "The Technical Paper Generator". What do you think?

—Needs Tenure

Dear Needs:
First, for the benefit of those readers who have not seen the system, allow me to outline it. The program takes a sentence as its input, rewrites the sentence four times to form an introduction (taking three of these separately to form an abstract), then makes up to five obvious statements about the sentence, expands each of these to a paragraph, and finally rewrites the introduction to form a conclusion. It also has a feature to perform a keyword search on recent publications and insert references.

Now certainly there are some merits to this sort of program, but the requirement of having to come up with a sentence in the first place could be a serious problem. I would recommend that you sell your program to a business school.

Dear Dr. Rob:
That was a cheap shot.

—Business School Graduate

Dear Grad:
No, no, no! Never say, "That was a cheap shot"! Always say, "The cultural benefits of the previously suggested implications are substantially negated by the sensationalistic slant inherent in their continuance of standard stereotypes."

Dear Dr. Rob:
Ok, smart guy! Which came first, the chicken or the egg?

—Brain Teaser

Dear Brain:
Neither. Aardvarks came first.

Dear Dr. Rob:
Clearly, there is some conflict between the frame representation of Minsky, on the one hand, and the "just new syntax" of Hayes on the other, possibly balanced by Winograd's declarative/procedural controversy, although I haven't heard much about that since Simons postulated "Are computers alive?", so how can this be reconciled with Etherington's inheritance of Brachman's "What is a IS-A?" in the light of McCorduck's "Machines who think" when Levesque and Nilsson cannot even agree over a lunch order and Steele lisps with Winston about Aristotle's rationality?

—Name Dropper

Dear Dropper:
Easily. Just peruse some of Gurevich's fuzzy logic.

Dear Dr. Rob:
Who is Number One?

—Vociferous Villager

Dear Vociferous:
You are Number Six. Number One is John Galt, or Ayn Rand, or Marshall McLuhan.

Dear Dr. Rob:

```
(setq question
      (and (prolog-hacks (friends me))
           (think (friends me)
                  (and (silly LISP)
                       (useless-for LISP AI))
                  :reason (proliferance-of '("(" ")")))))
(eval (think you question))
```

Parenthetically yours,
—Billy Bracket

Dear Billy:
I assert correct (friends) .

Dear Dr. Rob:
My speech recognizer has recently developed an amorous relationship with my vocal synthesizer. Their mutual affections are beginning to affect the research in the lab. What can I do?

—Confused

Dear Confused:
Machines have been falling in love for centuries. My car, for instance, has developed an affinity for the hoist at the local service station. You must try to persuade the pair to develop goals independently, inform them that this can only serve to strengthen their relationship by keeping it from turning inward and becoming destructive. Let them explore their feelings for others, possibly even terminals and printers. The speech recognizer could listen to poetry, while the synthesizer could perfect a foreign language. With understanding and respect from others in the lab, both should be able to return to full productive capacity.

Dear Dr. Rob:
I have heard that anthropomorphism is a serious problem for AI researchers. Is this true?

—Interested Interdisciplinarian

Dear Interested:
Not at all.

Dear "Dr. Rob":
Where do you get off with this "Dr." stuff? We've checked the records. You were probably lucky to get out of kindergarten. If you don't stop this nonsense we will have no choice but to tell the world about the time that you and (Unfortunately, there is not enough space to publish the remainder of this letter.—ed.)

(*Editor's note*: Please send your questions to Dr. Rob, c/o cscsi@noah.arc.cdn.)

NEXPERT OBJECT™

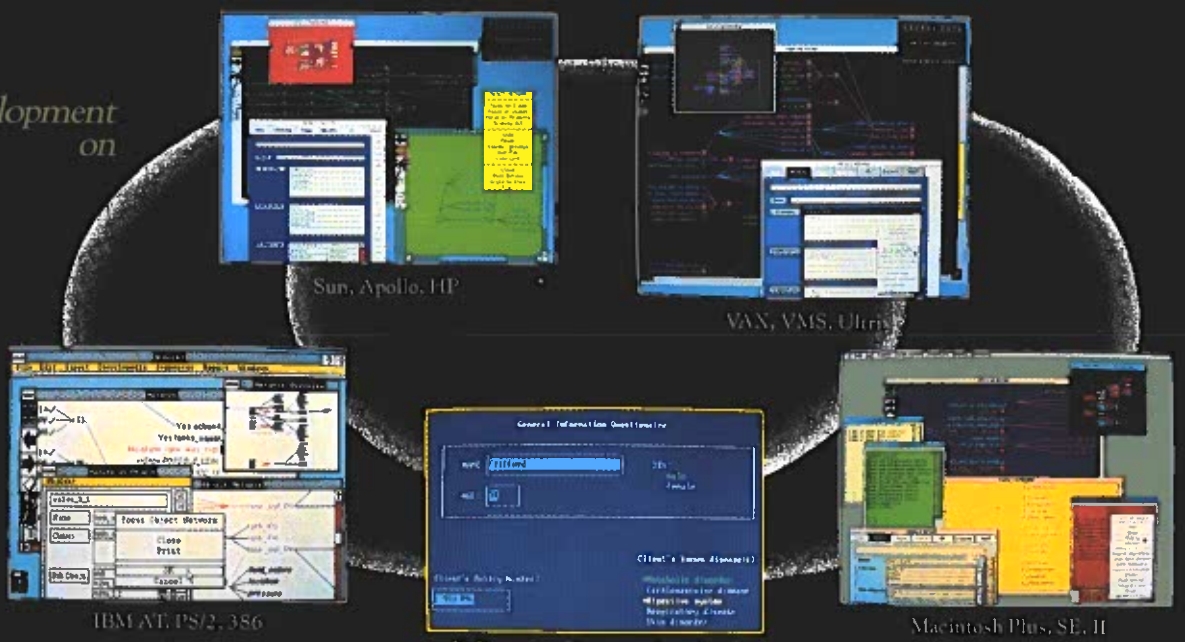

The Expert System Standard on Workstations

NEXPERT *OBJECT* is a hybrid rule- and object-based expert system tool with full cross-compatibility for development and delivery across all standard workstations. The Runtime AI Library in C provides the solution to embedded delivery systems in operational environments.

■ **EXPERT SYSTEM** □ HYBRID TOOL □ OBJECTS AND RULES □ WRITTEN IN "C" □ RUNTIME UNDER 300K BYTES ■ **RULES** □ PATTERN MATCHING □ INTEGRATED FORWARD AND BACKWARD CHAINING □ NON-MONOTONIC REASONING ■ **OBJECTS** □ CLASSES □ PROPERTIES □ META SLOTS □ MULTIPLE INHERITANCE □ METHODS □ DEMONS ■ **INTERFACE** □ GRAPHIC KNOWLEDGE EDITORS AND BROWSERS □ WINDOWS □ CUSTOM RUNTIME INTERFACE

DATAVIEWS is a dynamic graphics software tool kit. It allows users to develop object-oriented graphic interfaces for applications requiring displays which can be dynamically updated by NEXPERT. DATAVIEWS is sold by V.I. Corporation.

■ **OPEN ARCHITECTURE** □ INTEGRATION WITH EXTERNAL PROGRAMS □ LIBRARY AND STANDARD SYSTEM CALLS □ DIRECT ACCESS TO RELATIONAL DATABASES □ SQL DATABASE QUERIES. RDB, ORACLE □ DBIII,™ LOTUS 1,2,3™ ■ **HARDWARE** □ CROSS COMPATIBILITY □ IBM® AT, PS/2, 386 UNDER MICROSOFT WINDOWS,™ DOS, OS/2 □ MACINTOSH® □ $5,000 □ UNIX WORKSTATIONS UNDER X™ WINDOWS □ SUN, APOLLO, HP, VAX □ DEC® VAXSTATIONS II AND 2000 UNDER VMS/UIS □ $8,000 ■ **VARS, OEMS** □ TO QUALIFY FOR OUR VAR/OEM PROGRAM CONTACT NEURON DATA AT 415/321-4489 ■ **TRAINING** □ AVAILABLE THROUGH NEURON DATA, BECHTEL, DEC... ■ **DEC** □ A DIGITAL DISTRIBUTED SOFTWARE WITH DEC SUPPORT ■ **CLIENTS** □ LEADING CORPORATIONS USING NEXPERT FOR **DEVELOPMENT & DELIVERY** ■ BOEING □ NASA □ DEC □ IBM □ SHELL OIL □ E. I. DUPONT □ ARTHUR ANDERSEN □ HUGHES AIRCRAFT □ MCDONNELL DOUGLAS □ ROCKWELL □ G.T.E. □ AT&T □ POLAROID □ EDS □ GE □ LOCKHEED □ TRW □ U.S. ARMY □ U.S. AIR FORCE □ COMSAT □ NORTHERN TELECOM □ XEROX □ SIEMENS □ MERRILL LYNCH □ BOOZ ALLEN □ GM □ PROCTER & GAMBLE □ AETNA □ BECHTEL □ CHRYSLER □ BAYER □ HP... ■ **FOR MORE INFORMATION** □ CONTACT **NEURON DATA** □ 444 HIGH ST. □ PALO ALTO, CA 94301 □ 415/321-4488 □ ON-LINE BULLETIN BOARD

NEURON DATA

GIRICO

Groupe Interuniversitaire de Recherche en Informatique Cognitive des Organisations

Sponsored by :

CIPS Canadian Information Processing Society
ACFAS Association Canadienne Française pour l'Avancement des Sciences
AFCET L'Association Française pour la Cybernétique Economique et Technique
FIQ Fédération de l'informatique du Québec

Workshop chairman

Gérard Simian Univ. Laval (C)

Local Arrangement Chairman

Yves Hudon GIRICO (C)

Program Committee Chairman

Bernard Moulin Univ. Laval (C)

Program Committee

Lorne Bouchard	U. Q. à Montréal (C)
Michel Cayrol	Univ de Toulouse (F)
Eugène Chouraqui	GRTC, Marseille (F)
Pauline Côté	U.Q à Rimouski (C)
Philippe Duchastel	Université Laval (C)
Jean Louis Ermine	Bordeaux (F)
Henri Farreny	Univ. de Toulouse (F)
André Flory	Univ. Lyon3 (F)
Richard Fournier	Min. Education Québec (C)
Claude Frasson	Univ. de Montréal (C)
Car l H. Frederiksen	Univ. Mc Gill (C)
Gilles Gauthier	U. Q. à Montréal (C)
Jean F. Hue	Univ. Nantes (F)
Gilles Imbeau	U.Q. à Chicoutimi (C)
Martin Janta-Polczynski	Univ. Mc Gill (C)
Jacques H. Jayez	Univ. Nantes (F)
Jacques Kouloumdjian	Univ. Lyon1 (F)
Luc Lamarche	Univ. de Montréal (C)
Serge Larochelle	Univ. de Montréal (C)
Rolland Lesuisse	Institut Inf. Namur (B)
Ruddy Lelouche	Univ. Laval (C)
Georges H. Moll	Univ. Lyon1 (F)
Joel Muzard	CCRIT, Montréal (C)
Cao Lieu Nguyen	U.Q. à Montréal (C)
Pierre Plante	U.Q. à Montréal (C)
Spencer Star	Univ. Laval (C)
Richard St Denis	Univ. de Sherbrooke (C)
André Théoret	Univ. de Sherbrooke (C)
Jean Vaucher	Univ. de Montréal (C)

ico '89

May 29 to June 1,1989 Québec city, Canada

INTERNATIONAL WORKSHOP ON COGNITIVE INFORMATICS APPLIED TO ORGANIZATIONS

Impact of artificial intelligence and cognitive sciences in organizations in the nineties

Call for papers

Original papers are being sought in all areas related to applications of artificial intelligence (AI) and cognitive science (CS) in organizations. Papers concerning basic research, tools and techniques should focus on one or more aspects of: representation, acquisition, processing and communication of knowledge; machine learning; problem solving; natural language processing; person / machine interfaces; impact of AI and CS in organizations, knowledge-based system design methods, intelligent tutoring systems; applications of AI and CS in industry; expert systems.

author's submission:	**August 30, 1988**
notification of acceptance:	**October 30, 1988**
camera-ready final papers:	**December 10, 1988**

Instructions to authors

Authors are invited to submit four copies of their papers, not exceeding 4000 words, double spaced, in French or English. Submitted papers should include a page with : title of paper, author's name and address, phone number, a 10 lines abstract, a list of keywords. Submitted papers should reach the program committee chairman before August 30 at the following address :

The workshop will include:
tutorials (May 29), industrial conferences (May 30), technical conferences (May 31 and June 1).
An exhibition will be held concurrently with the worshop. Commercial products (hardware and software) and research prototypes will be exhibited.

For information : Prof. Bernard Moulin
Colloque ICO'89
Université Laval dépt. d'Informatique
Ste Foy, Québec G1K 7P4 Canada
phone: (418) 656-5580 or 656-7979

Short Takes

CSFGR Address Update

The president of the Canadian Society for Fifth Generation Research, Dr. Nick Cercone, is now the director of the Centre for Systems Science at Simon Fraser University. Henceforth, the address for CSFGR/SCRSCG is now: c/o Dr. N. Cercone, CSS, Simon Fraser University, Burnaby, British Columbia V5A 1S6.

CIAR Moves

The Canadian Institute for Advanced Research, the non-profit "think-tank" for determining long-term research important to Canada's future, has moved to: 179 John Street, Suite 701, Toronto, Ontario M5T 1X4. Phone: (416) 971-4253.

Sony and Symbolics Announce Strategic Alliance

New York, May 3, 1988 — Sony Microsystems Company, based in Palo Alto, Ca., and Symbolics Inc., based in Cambridge, Ma., have announced an alliance in which Symbolics will act as a reseller and third party service organization for Sony's NEWS family of UNIX™ workstations.

Symbolics will initially incorporate the Sony NEWS NWS-711 and NWS-841 workstation into its product offering. Symbolics' customers can use NEWS as general-purpose workstations for the development and delivery of UNIX applications. NEWS and Symbolics systems may be connected in a network with the NEWS stations accessing the Symbolics machines as "knowledge servers".

As part of the alliance, Sony and Symbolics will share new product plans and engage in technical and sales training, and cooperative marketing programs. In the future, Symbolics plans to embed its Ivory™ single-chip symbolic processor in the general purpose workstations, and provide a version of Cloe™ (Common Lisp Operating Environment).

Symbolics Inc. is a leading developer, manufacturer, and marketer of symbolic processing systems, with revenues of US$103 million for the year ended June 30, 1987. Sony Microsystems Company is a division of Sony Corporation of America, a subsidiary of the US$10.2 billion Sony Corporation, Tokyo, Japan. Sony Microsystems was formed in November 1987 to market state of the art computer hardware and software manufactured by Sony in the United States.

Federal Scholarship Program

As part of the Cdn$1.3 billion in science and technology funding committed by the federal government over the next five years, a major undergraduate scholarship program has been established. The first year of the program, beginning in the Fall of 1988, will award 2,500

Prise de vue

SCRSCG déménage

Le président de la Société canadienne pour la recherche sur les systèmes de cinquième génération, le Dr. Nick Cercone, a déménagé pour prendre la direction du Centre for Systems Science à l'Université Simon Fraser. L'adresse de la SCRSCG/CSFGR est maintenant: a/s Dr. N. Cercone, CSS, Simon Fraser University, Burnaby, Colombie Britanique V5A 1S6.

L'ICRA déménage

L'Institut canadien pour la recherche avancée, le groupe d'experts sans but lucratif qui a pour objectif de déterminer quels types de recherche ont une importance a long terme pour l'avenir du Canada, a déménagé au 179 John Street, Suite 701, Toronto, Ontario M5T 1X4. Phone: (416) 971-4253.

Sony et Symbolics annoncent une alliance stratégique

New York, le 3 mai 1988 — Les compagnies Sony Microsystems de Palo Alto, Ca., et Symbolics Inc., de Cambridge, Ma., ont annoncé une entente qui fait de Symbolics le distributeur ainsi que la structure responsable de l'entretien aux clients des postes de travail NEWS utilisant UNIX™ de Sony.

Symbolics offirra d'abord les modèles NEWS NWS-711 et NWS-841. Les clients peuvent utiliser ces postes de travail pour développer des applications UNIX. Les postes de travail NEWS peuvent être inclus dans un réseau de machines Symbolics où ces dernières agissent en tant que serveurs pour la mise en commun des connaissances.

Cette entente stipule que Sony et Symbolics partageront la conception de nouveaux produits et mettront en commun leurs programmes de formation technique et de vente ainsi que leurs programmes de mise en marché. A l'avenir, Symbolics prévoit utiliser son processeur Ivory™ et une version de Cloe™ (Common Lisp Operating Environment) dans ces postes de travail.

Avec 103 millions dollars americain, pour l'année se terminant le 30 juin 1987, Symbolics est un plus important manufacturier de systèmes de traitement symbolique. La compagnie Sony Microsystems est une division de la Sony Corporation d'Amérique, elle-même une filiale du géant Sony Corporation de Tokyo (10.2 milliards dollars). Cette compagnie a été formée en 1987 pour la mise en marché des dernières découvertes de Sony d'Amérique en matière de logiciel et de matériel.

Programme fédéral de bourses

Une partie des 1.3 milliards dollars canadien, alloués aux sciences et à la technologie, par le gouvernement fédéral, pour les cinq prochaines années, sera utilisée pour un vaste programme de bourses d'études sous-diplômées. Le

Canada Scholarships worth Cdn$2,000 each. The scholarships will be renewable up to a further three years.

The scholarships will be allocated to post-secondary institutions based on the number of undergraduate degrees awarded in the eligible disciplines. The eligible fields and disciplines include agriculture and biological sciences, engineering and applied sciences, and mathematical and physical sciences.

A minimum of 1,250 Canada Scholarships will be awarded to women. Particular emphasis will be placed on nominating women who are entering disciplines where their representation is lowest. Any Canadian citizen or permanent resident, registered as a full-time student in an eligible discipline in a recognized post-secondary institution, is eligible for a Canada Scholarship. Students are required to submit an application to the institution of their choice. For the first year of the program, institutions are free to nominate scholars of their choice in any of the eligible disciplines.

PRECARN RFP Update

Earlier this year, PRECARN, the Precompetitive Applied Research Network posted an RFP for research and development activities in artificial intelligence and robotics (see *Canadian Artificial Intelligence*, April 1988).

PRECARN received 28 proposals for feasibility studies by the deadline of April 18. These proposals were reviewed by a technical advisory committee that in turn made recommendations to the Board of PRECARN. At the time of press, the committee had approved 7 proposals unconditionally for recommendation. Another 4 were recommended to the Board subject to certain conditions being met. Details will appear in the October issue of this magazine.

For more information contact: PRECARN Associates Inc., 30 Colonnade Rd., Suite 300, Nepean, Ontario K2E 7J6. Phone: (613) 727-9576.

Canada to Participate in International Space Station

The federal government has approved Canada's participation in the U.S.-led international Space Station, the major initiative within Canada's dynamic Space Program. Canada's investment to participate in the Space Station Project is estimated to be just under Cdn$1.2 billion over 17 years.

This cost covers several aspects of the Space Station Project, including the design and construction of the Mobile Servicing System (MSS), the development of the strategic technologies necessary for the project, the operation costs associated with the MSS, Canada's share of the common systems operating costs of the entire Space Station, and the User Development Program.

The MSS will incorporate technologies which are strategically important for Canada. The Strategic Technology Automation and Robotics (STEAR) program has been established to help industries to exploit industrial applications of these MSS technologies. Under STEAR, firms, universities and research organizations will be contracted to carry out work which will complement the technology development required to build MSS.

To date, two RFPs have been tendered by the Department of Supply and Services, relating directly to

programme, qui commence à l'automne 1988, distribuera 2500 bourses Canada de 2000$ chaque. Les bourses seront renouvelables pour un maximum de trois ans.

Les bourses seront réparties aux institutions post-secondaires selon leur nombre annuel de diplômés dans les matières éligibles. Ces dernières comprennent les sciences de l'agriculture et de la biologie, le génie et les sciences appliquées, les sciences physiques et les mathématiques.

Au moins 1250 bourses iront à des femmes, en particulier dans les domaines où leur nombre est le plus faible. Tout citoyen canadien ou résident permanent inscrit à temps complet dans l'une des matières éligibles à l'une des institutions post-secondaires reconnues, est éligible à une bourse Canada. Les étudiants doivent soumettre leur demande à l'institution de leur choix. Pour la première année du programme, les institutions choisiront elles-mêmes leurs récipiendaires.

Programme de subventions de PRECARN

Plus tôt cette année, PRECARN, le réseau préconcurrentiel de recherche appliquée, a annoncé un programme de subventions pour les activités de recherche et développement en intelligence artificielle et en robotique (voir *IA au Canada*, avril 1988).

PRECARN a reçu 28 propositions pour des études préliminaires avant l'échéance du 18 avril. Ces propositions ont été évaluées par un comité consultatif technique qui a à son tour fait des recommendations au conseil de PRECARN. A l'heure de tombée, le comité avait approuvé 7 propositions sans conditions pour recommendation. Quatre autres furent recommandées au conseil sujet à certaines conditions. Les détails apparaîtront dans l'édition d'octobre de ce magazine.

Pour de plus amples informations, contacter: PRECARN Associates Inc., 30 Colonnade Rd., Suite 300, Nepean, Ontario K2E 7J6. Phone: (613) 727-9576.

Participation canadienne à la station spaciale internationale

Le gouvernement fédéral a approuvé la participation canadienne au projet de station spaciale internationale mené par les Etats-Unis. Cette participation constitue la plus importante partie du programme spatial canadien. L'investissement sera de 1.2 milliards $Cdn répartis sur 17 ans.

Ce coût comprend, entre autres, la conception et la construction d'un système mobile d'entretien (le MSS), le développement des technologies stratégiques requises par le projet, les frais de service associés au MSS, la part du Canada dans les frais communs, et le programme de développement des usagers.

Le MSS requérera des technologies qui sont d'importance stratégique pour le Canada. Le programme STEAR (Strategic Technology Automation and Robotics) a été créé afin de faire profiter l'industrie des applications découlant de la technologie du MSS. Compagnies, universités et organisations de recherche recevront des contrats pour compléter la technologie du MSS.

Jusqu'à date, le ministère des approvisionnements et services a établi deux projets touchant directement à l'IA. Le premier porte sur l'automatisation de tâches de soutien tel que l'entraînement des équipes, la gestion des

artificial intelligence. The first, "Automation of Operations", includes the automation of such support operations as operations planning; crew training; resource management; data acquisition, analysis and management; system health monitoring and failure anticipation; technical information management; anomaly/exception handling; and fault detection and diagnosis. The second RFP pertains to "Development of a Knowledge-based Health Monitoring and Planning System for Space Station MSS Power Management-Phase I". Strong responses were received to both RFPs. It is anticipated that before the end of June, up to 6 contracts will be awarded for work on the first RFP. A further contract will be awarded at a later date for the second RFP.

New Products

New Release of Symbolics Operating System

Symbolics Inc. has announced an extensive upgrade of its Genera® operating system software, used by all Symbolics computers. Genera 7.2 features higher performance and the user interface has been extended to be consistent across system components. There are eight new manuals in the documentation set, five of them are specifically for new users. The company began shipping the system in mid-April.

NRC Transfers Harmony Software to Canadian Distributor

P-CAN Research Inc., a company specializing in high performance automation and computing products, will commercialize a sophisticated software system called Harmony, developed by the National Research Council. Harmony is a distributed operating system which allows different machines to communicate and work together in real-time.

Developed by NRC's Division of Electrical Engineering for real-time control of industrial and automation equipment, Harmony has a wide range of potential applications in industry, communications, and aerospace command and control. For more information, contact: Fred Joneidi, P-CAN Research Inc. Phone: (416) 674-6600.

Sun enters PC Market

Sun Microsystems has announced its entry into the PC marketplace with the introduction of the 32-bit Sun 386i series of computers. The system allows execution of off-the-shelf DOS applications within the traditional Sun Unix™ environment. A user-friendly, icon-based, multi-tasking interface has been developed. Files may be shared by both UNIX and DOS applications.

Machines are available in a variety of configurations, delivering between 3 and 5 MIPS. The base configuration, the Sun 386i/150 monochrome workstation, which comes with a 15 inch monitor, 4MB memory, keyboard, mouse, and SunOS, is priced at Cdn$13,600. This system runs as a diskless node in a networked environment. The 386i/250 colour system,

ressources, l'acquisition, l'analyse et la gestion des données, le contrôle du système et l'anticipation des pannes, la gestion de l'information technique, le traitement d'anomalies et d'exceptions, la détection et l'analyse d'erreurs. Le second projet vise à développer un système de contrôle et de planification à base de connaissances pour la gestion de l'énergie du MSS. Les deux projets ont entraîné plusieurs réponses et l'on s'attend à allouer 6 contrats pour le premier projet avant la fin de juin. Un contrat pour le deuxième projet viendra plus tard.

Nouveaux produits

Nouvelle version du système d'exploitation de Symbolics

Symbolics Inc. a annoncé une version substantiellement améliorée de son logiciel de système d'exploitation Genera®, utilisé par tous les ordinateurs Symbolics. Genera 7.2 donne une plus haute performance et l'interface avec l'usager a été étendu de façon à être uniforme à travers les différentes composantes du système. L'ensemble de documentation comprends huit nouveaux manuels, dont cinq s'adressent spécifiquement aux nouveaux usagers. La compagnie a commencé à expédier le système à la mi-avril.

Le CNR transfère le logiciel Harmony à un distributeur canadien

P-CAN Research Inc., une compagnie spécialisée dans les produits informatiques et d'automatisation de haute performance, commercialisera un logiciel sophistiqué, nommé Harmony, développé par le Conseil national de recherche. Harmony est un système d'exploitation réparti qui permet à des machines différentes de communiquer et de travailler ensemble en temps réel.

Développé par la Division de génie électrique du CNR pour le contrôle en temps réel de l'équipement industriel et d'automatisation, Harmony a un domaine d'applications potentielles important dans l'industrie, en communication et commutation, et commandement et contrôle aérospatial. Pour plus d'informations, contacter: Fred Joneidi, P-CAN Research Inc. Phone: (416) 674-6600.

Sun fait son entrée sur le marché des PC

Sun Microsystems a annoncé son entrée sur le marché des ordinateurs personnels avec l'introduction de la série d'ordinateurs à 32 bits Sun 386i. Le système permet l'exécution d'applications DOS à l'intérieur de l'environnement Unix™ traditionel de Sun. Un interface multitâche amical à base d'icônes a été développé. Les fichiers peuvent être partagés par les applications UNIX et DOS.

Diverses configurations produisant entre 3 et 5 MIPS sont disponibles. La configuration de base, le poste de travail Sun 386i/150 monochrome, dispose d'un écran de 15 pouces, 4MB de mémoire, un clavier, une souris, et SunOS, et se vend pour 13 600$ canadien. Ce système opère comme un noeud sans disque dans un réseau. Le

which includes a 14 inch monitor and a 91 MB hard disk, is priced at Cdn$18,700.

GoldWorks on Sun

Gold Hill Computers has announced plans to port their GoldWorks expert system building tool to Sun Microsystems' new Sun 386i™ family of workstations. GoldWorks will run on top of Sun Common Lisp and will be shipped in the fourth quarter of 1988. The software will be priced the same as the PC version of GoldWorks: US$7,500 for one copy, US$4,995 for quantities of three or more.

XEROX 1100 Series Farewell

XEROX will stop manufacturing the 1100 Series Lisp machine and its descendants this summer. The company has a US$200 million OEM agreement with Sun Microsystems and is porting its InterLisp environment to the Sun machines. The combined environment should be available at the end of this summer.

Unix™ on Mac II

A/UX, Apple's version of the UNIX operating system is available for the Mac II. The system began shipping in mid-April and is a fairly complete version of AT&T's System V 2.2 operating system. Existing Macintosh applications may be run under the UNIX environment. The operating system and an 80MB hard disk is priced at Cdn$5,299 for an external hard drive, and Cdn$5,137 for an internal drive. Manuals are a separate product and cost Cdn$1,048.

système couleur 386i/250, inclus un moniteur de 14 pouces, un disque dur de 91 MB, et se vend pour 18 700$.

GoldWorks sur Sun

Gold Hill Computers a annoncé son intention de transporter son outil de développement de systèmes experts GoldWorks sur la nouvelle famille de postes de travail Sun 386i de Sun Microsystems. GoldWorks opèrera à l'intérieur de Sun Common Lisp et sera livré dans le quatrième trimestre de 1988. Le logiciel se vendra au même prix que le version PC de GoldWorks: 7500$ americain pour une copie, 4995$ pour des quantités de trois ou plus.

La série 1100 de XEROX fait ses adieux

XEROX cessera de fabriquer la série 1100 de machines Lisp cet été. La compagnie a une entente d'OEM de 200 millions dollars americain avec Sun Microsystems et est en train de transporter son environnement InterLisp sur les machines Sun. L'environnement combiné devrait être disponible cet été.

Unix™ sur le Mac II

A/UX, la version Apple du système d'exploitation UNIX, est disponible pour le MAC II depuis la mi-avril. Ce système est une version relativement complète du système d'exploitation V 2.2 de AT&T. Les logiciels existant déjà pour le Mac peuvent rouler sur A/UX. Le système d'exploitation et un disque rigide de 80MB coûtent 5299$ canadien pour un disque externe, et 5137$ pour un disque interne. Les manuels sont vendus séparément à 1048$.

A Proposal for the Creation of SIGET

by Philippe Duchastel

It would appear that the time is ripe for the creation of a mechanism which would permit Canadian researchers interested in the application of artificial intelligence to education and training to better know one another and to promote their field of research. Indeed, Canadian activity in this field is both expanding and becoming more visible, as indicated by the following signs:

- creation in 1987 of an AI in Education sub-committee of the Associate Committee on Instructional Technology (ACIT) of the National Research Council of Canada
- initial listing of Canadian "AI in Education" projects, compiled by Jim Greer of the University of Saskatchewan
- hosting of an international Intelligent Tutoring Systems conference by the University of Montreal group in June 1988
- initial planning of an AI in Education Summer Institute for 1989

All of these signs show that interest in this field is growing and that coordination and promotion will serve the interests of both the practitioners and the eventual users of this technology. Some form of association would seem desirable to achieve these aims, hence this proposal for the creation of a Special Interest Group on Education and Training (SIGET) within CSCSI.

The goals of this new SIG would be along the following lines:

- continuing to sustain interest in this field within the Canadian AI community (a long history of interest is already established within CSCSI)
- promoting interest in this field within the Canadian education and training communities (interest is beginning to grow but is still unfocused)
- offering a mechanism of identification and communication for Canadian researchers active in this field (for example, through articles and news notes in *Canadian AI*, through conferences and institutes)
- coordination with other bodies sharing interest in this field (e.g. ACIT, AAAI, NATO)
- leadership in the creation of an international association specifically devoted to this field (coordination is currently sorely lacking, there being yet no formal American nor European association in this field).

The above list of SIG functions is of course merely a

Philippe Duchastel was on the program committee for the International Conference on Intelligence Tutoring Systems, in Montreal in June.

proposal and will be polished by the SIG members.

If you believe you would be interested in becoming a member of SIGET (irrespective of whether you are currently a CSCSI member or not), please indicate this interest by communicating with myself at the following address:

Dr. Philippe Duchastel, director, Laboratoire d'Intelligence Artificielle en Education, Bureau 1466, Pavillon Charles-De-Koninck, Université Laval, Québec, Canada G1K 7P4. Phone: (418) 656-5085 or 656-3769. Net North: 1160001 @ LAVALVX1

Une proposition pour la formation de GISEF

par Philippe Duchastel

Comme l'indique la participation québécoise et canadienne à la Conférence internationale sur les systèmes tutoriels intelligents organisée par nos collègues du Groupe Héron de l'Université de Montréal, l'intérêt au Canada francophone pour l'intelligence artificielle appliquée à l'éducation ou à la formation croît rapidement (de fait, il n'existait guère il y a trois ans).

Par contre, aucune coordination n'existe présentement dans ce domaine à aucun niveau: ni québécois, ni canadien, ni international. De plus, alors que notre domaine spécialisé de recherche est bien compris par la communauté de recherche en intelligence artificielle (le premier système EIAO remonte en effet à 1970), il n'en va pas du tout de même dans la communauté éducative ou centrée sur la formation. Deux efforts sont donc requis aujourd'hui pour mieux servir les intérêts des chercheurs oeuvrant en EIAO: coordination et promotion. Il en découle la souhaitabilité d'un mécanisme organisationel tel que celui proposé ici.

Aujourd'hui au Québec, deux organismes remplissent un certain rôle de coordination (il s'agit du CRIM et du GIRICO), mais de façon encore très modeste. Au Canada, le Comité associé du CNRC en Technologie pédagogique joue également un rôle de leadership dans ce domaine.

L'avantage du Groupe d'Interêt Specialisé en Education et Formation (GISEF) comme groupe affilié à la Société Canadienne pour l'Etude de l'Intelligence par Ordinateur (SCEIO) réside principalement dans la nature pan-canadienne de cet organisme (d'où une certaine masse critique initialle) et également dans l'existence d'un excellent magazine bilingue (*Intelligence Artificielle au Canada*) qui peut servir les besoins de communication des chercheurs en EIAO.

Il y a lieu de préciser l'étendue des intérêts représentés en principe par le GISEF. De façon générale, c'est l'EIAO: l'Enseignement Intelligemment Assisté par Ordinateur, tel qu'appliqué soit en éducation formelle ou informelle, soit en formation. Il n'y a pas lieu d'être restrictif dans la définition du domaine (par exemple, les systèmes experts en planification pédagogique, ou le design d'environnements d'apprentissage intelligents). Par ailleurs, il y a lieu d'exclure tout ce qui ne comprend pas implicitement une approche à base d'intelligence artificielle (par exemple, la plupart des recherches sur LOGO). Ceci n'est que strictement indicatif puisque les fonctions et la nature du GISEF seraient éventuellement précisées par les chercheurs et chercheures qui le composeront.

Si donc le GISEF peut vous intéresser (que vous soyez ou non membre du SCEIO), je vous prie de me l'indiquer en communiquant avec moi:

Dr. Philippe Duchastel, directeur, Laboratoire d'Intelligence Artificielle en Education, Bureau 1466, Pavillon Charles-De-Koninck, Université Laval, Québec, Canada G1K 7P4. Phone: (418) 656-5085 ou 656-3769. Net North: 1160001 @ LAVALVX1

AI and Resource Industries

by Connie Bryson

L'IA et l'industrie des ressources naturelles

RÉSUMÉ: Même si les prises de positions et les rapports de consultants ont régulièrement exposé le potentiel des applications de l'IA dans l'industrie des ressources naturelles au Canada, ce n'est que récemment que l'industrie et les compagnies de logiciel ont pris conscience de ce potentiel. Les activités des dernières années ont beaucoup rétréci l'écart qui séparait les voeux exprimés des produits disponibles sur le marché. Cet écart existe toujours, mais plusieurs experts prévoient une révolution dans les applications de l'IA dans des industries comme celles du pétrole et du gaz, des mines, de la forêt, des pêches, et de l'agriculture.

Although position papers and consultants' reports have regularly expounded the virtues of AI applications in Canada's resource industries, only recently have industry and software companies acknowledged that there are commercial opportunities for AI in the resource sector. The activity in the last few years has done much to narrow the gulf between what's being said and what's on the market.

But a gap still exists and it could be seen quite clearly at the 90th Annual General Meeting of the Canadian Institute of Mining and Metallurgy held in Edmonton in early May. The theme for the meeting was "The 1990's — Formulas for Success".

Few participants would dispute the assertion that the use of advanced computer technology is one of these success formulas. However the trade show at the conference told a different story. It featured many large machines and very little computer technology, hardly any of which could be classified as AI. It seemed the technical sessions and trade show were worlds apart.

But don't despair, say many analysts. The revolution really is coming.

The Times They are A-Changin'

"Historically, Canada has emphasized production and imported the technology," says Jeff Pallister, Vice-President, Pallister Resource Management Ltd., Calgary.

"We could continue to import. Advanced technology is coming from the U.S., Japan, Scandanavia. But we could also develop the technology here — that's the opportunity I see. After all, Canada already has the resource industry expertise."

One of the obstacles in the way of developing AI technology for the resource industries is the very poor linkages between software companies and resource companies. "They operate in different environments and don't communicate," Pallister says.

"But there's definitely a market. By and large, resource companies aren't developers; they will buy technology if it exists. However suppliers are often small and don't have the money for development. It's difficult to get going."

Difficult is an understatement, according to Robert Czinner, President, Integrated Wood Research Inc. He is trying to attract venture capital for his company's advanced technology lumber remanufacturing process.

"The bottom line in Canada is that advanced manufacturing is paid lip service," he says. "They (the financial community) love to talk about it until the question of their participation comes up. It's a fundamental problem in Canada. We're going to miss the boat."

Czinner says that advanced technology cannot be assessed using conventional cost accounting formulas. He says new applications require a different kind of evaluation — and a willingness to take risks.

"People think that if they keep asking technical questions the risk will disappear. But the risk can't disappear if no one has done it before."

Despite his frustration over funding, Czinner remains convinced advanced technology is the best way to efficiently use natural resources and he is optimistic the merits of his system will eventually turn doubters into believers.

Grant Thomas, President of Ottawa's CAIP Corporation, says two factors govern the "climate" influencing AI applications in the resource industries. (Thomas is talking specifically about expert systems.)

"It's only now that we have some people in Canada — not enough I may add — with the expertise to build expert systems. Two or three years ago, we didn't. But now our capability base is maturing," he says.

"At the same time there is an increasing awareness of AI in the user community. Companies are starting to think that maybe there's something in this. Large companies have a couple of people who are starting to experiment with AI. A consumer pull is developing."

Organizations such as Gordon MacNabb's PRECARN Associates Inc. are stimulating the demand/supply convergence. This consortium of Canadian companies has pledged financial and management help for AI research. Thomas calls PRECARN's work "absolutely fundamental" to the transfer of AI technology from research to application.

MacNabb reports that many of the proposals submitted to PRECARN for its first call for proposals relate directly to the resource industries.

"Although Canada lags behind other countries, we have a sprinkling of exceptional talent across the

Connie Bryson is a free-lance technical writer based in Vegreville, Alberta.

country," MacNabb says. "At PRECARN, we're trying to bring academia and industry together. I'm getting a very positive message from government that the money will be available for these projects."

Apart from direct funding, the government is encouraging AI technology transfer in another way. The National Research Council (NRC) has embarked on a focused program in technologies related to the resource industries — robotics, biotechnology and computer process control. An internal NRC steering committee was set up to consider how NRC laboratory programs can be more relevant to the resource industries — agriculture, fisheries, forestry and minerals.

The committee produced a preliminary document in 1987 and is now in the process of testing whether the preliminary ideas are correct, says Dr. Phil Cockshutt, NRC's Executive Director, Engineering Programs. Workshops are planned for the mining and forestry sectors in the fall of this year.

"I don't believe there's any difference between applying AI to natural resources or to manufacturing," says Peter MacKinnon, Cognos Inc. "It's just that the resource industry is only now starting to catch on."

The Importance of the Resource Sector

Besides the inherent applicability of many AI technologies to the resource industries, the government sees another reason to push AI in these industries — the fact that the resource industries are of enormous economic significance to Canada. They are directly responsible for about 14 per cent of the annual Gross Domestic Product and for almost 35 per cent of gross exports.

But numbers do not tell the whole story. In the Canadian context, there is a strong geographic element to be considered — resource industries constitute the major source of economic activity outside the three or four major urban centres in the country.

As important as it is now to Canada's economy, the future importance of the resource sector is uncertain. Markets are saturated with many primary commodities while lists of alternative materials grow. Canada has many new competitors with rich resources, cheap labour and transportation.

"For example, the mining industry must now compete with Indonesia with its rich ores and cheap labour," says PRECARN's MacNabb. "It's essential that the resource industries take advantage of AI. It's the only way they'll keep competitive."

Getting the Ball Rolling

Within a company, the desire to use AI technology does not develop overnight. It's usually a matter of convincing top administration that the work is not an esoteric investigation but one with practical applications.

"It often starts with someone internally who's interested in AI," says Cognos' MacKinnon. "Or sometimes a contract R&D firm makes a pitch. Other times the government, through NRC's IRAP program for example, plays a seed role."

Joel Ouriou, Technology Planning, Esso Resources, has found that as the need for specialized hardware for AI work disappears, the interest in AI systems at Esso grows. The company sees applications in accounting and marketing as well as in science and engineering.

"We're now looking at AI as we do any other technology," Ouriou says. "Its value will come when it can be integrated with our existing technology and not as a specialized stand-alone technology."

Ouriou says Esso is "at the stage of trying to understand the technology. We don't have a lot of skills inside the company now but over time we will build them up.

"I look at AI as another way of delivering software. It is a tool that will be used where it's appropriate. We're identifying those areas now."

Dr. Bev Smith, from Victoria-based Acquired Intelligence Inc., says the AI consulting business is still a matter of knocking on industry's door rather than having industry show up at the consultant's doorstep.

"But the situation is improving. More and more companies are looking into AI," she says.

"It's not a case anymore of industry hanging on to see if the hype is real. They've seen the savings and they are real. This isn't leading-edge risky technology anymore. We can't automate everything but if applications are chosen carefully, there is enormous potential."

AI Applications

According to Jeff Pallister, Pallister Resource Management, as far as adopting AI technology goes, the oil and gas industry is the most advanced sector of the traditional resource industries. Mining and forestry are somewhere in the middle, with fisheries and agriculture bringing up the rear.

Pallister has prepared an outline of expert system applications in the oil and gas industry. It includes nearly all upstream oil and gas activities:

- Geology
  - — basin evaluation
  - — stratigraphic correlation
- Geophysics
  - — surveying
  - — interpreting seismic data
  - — feature identification
- Drilling
  - — mud formulations
  - — interpretation of well logs
  - — fracture program planning
- Support
  - — evaluation and bidding
  - — economic analysis.

Pallister speculates that part of the reason oil and gas technology is more advanced than other sectors stems from the nature of the industry itself. The industry is concentrated in head offices in Alberta whereas farming and fishing are fragmented industries with end users far removed from R&D activity. And in agriculture, the poor farm capital picture has impeded developments in technology.

But this is not to say that farming and fishing are unsuited to the implementation of AI technology. For example, a study by Econome Consultants Inc. identified areas in Canada's fisheries where technological innovation could double the total wealth generated by fisheries over the next 20-30 years.

Pallister lists the following potential applications of expert systems in the agri-food industry:

- pest management programs
- troubleshooting mechanical systems
- diagnosis of animal health problems.

For forestry, he lists:

- interpretation of remote sensing data for resource

evaluation and management
- forest fire management systems
- manufacturing engineered components.

Raman Janakiraman from the Advanced Computing Technology Centre, University of Calgary, prepared this summary of AI application areas in the minerals industry:
- geo-sensing
- rock fragmentation
- intelligent mining systems
- diagnosis/evaluation
- process control and planning.

The following section contains summaries of some of the AI systems either under development or in use in Canada. It is intended as an indication of activity and does not constitute a complete listing.

Oil and Gas

When discussing expert system applications, many sources still refer to PROSPECTOR, a consulting geology program developed by the Stanford Research Institute and the U.S. Geological Survey about 10 years ago.

However there are a number of other, more recent examples of successful expert systems. This is particularly true in the oil and gas industry where large amounts of data are obtained and processed in every phase of operations.

One such expert system is Logmate Assistant, a new product from Calgary-based D&S Knowledge Systems Inc. It was developed jointly with the Alberta Research Council.

Logmate Assistant analyzes well logs to determine what is actually down an oil or gas well. It allows junior employees to perform the complex well log analyses that are usually done by experienced geologists. Logmate Assistant is the first commercially available system of its kind and one of the first expert systems that can be operated on a PC.

Gerald Parks, D&S business development manager, says the impetus for developing the system came from the fact that "a lot of expertise is starting to retire. This (an expert system) is a way to hold on to people."

Parks says the company's most difficult decision during Logmate development centred around choosing the platform. "We were unsure whether it should be on an inference engine like a Symbolics machine or a PC," he said.

"You can get into an awkward mess with hardware. With a bigger machine, many people know it and know what software works on it, but with a PC, you know the market is there."

Mining

Jack Scrimgeour, a senior advisor at NRC, reports that several countries, including France, Canada, Great Britain, Australia and the U.S., have done studies on robotic applications in mining.

"Although these were separate reports, they all came to approximately the same conclusion — that there's not much likelihood of using normal industrial robots in mining," he says. "They're just not suited to the mining environment although they may be used in ancillary roles like fixing trucks."

What Scrimgeour does expect to see is the "continual development of on-board control systems in mining equipment.

"Much of the mining equipment currently in use has on-board microprocessors to assist the operator in guidance and control. The next step is having the operator in a location away from the cutting face. Some day — and even mining people can see this now — we may have the operator on the surface."

Forestry

Just as in mining, automation is also underway in forestry. Dr. Peter Lawrence, Electrical Engineering, University of British Columbia, is heading a project on the control of machines in rugged terrain. The work incorporates vision, diagnostics and AI simulation.

Working in real-time, the prototype vision system monitors the joint angles on the boom and stick of a Caterpillar excavator. Using this system, the operator can run the machine from a remote location. The machinery is therefore easier and safer to use and can be run more efficiently.

"AI technology is fairly fragile. Forestry isn't like manufacturing where the environment is clean, controlled, and has defined boundaries," explains Dr. Alan Mackworth, a UBC computer scientist who is collaborating with Lawrence on the project. "Rather than trying to see the trees and all obstacles in the environment, we decided to look at the excavator arm.

"We can't have the system do everything. The key is to look for niche problems."

Expert systems are one of the advanced technology components that make up the lumber re-manufacturing process developed by Robert Czinner's company, Integrated Wood Research Inc. The Toronto company's proprietary technique takes non-commodity grade lumber and creates certain dimensions of commodity-grade lumber that are in short supply from primary mills.

Fisheries

Under a contract from the British Columbia Department of Fisheries, CAIP Corp. has just finished work on an expert system prototype now being tested in B.C.

When a forestry company decides to log an area, the logging plan must first be evaluated by fisheries officers to determine what impact the disturbance will have on fish habitat. Instead of combing through the B.C. Fish Habitat Guidelines book, an officer enters data to describe the logging operation. The system, which contains all the rules and regulations, prints out its recommendation along with the relevant sections of the guidelines.

Several of the large fish processing companies have expert systems in development. These systems use oceanographic and other types of data to predict where fish are likely to be and then integrate this information with market demand statistics.

Agriculture

Victoria-based Acquired Intelligence Inc. is working on two agriculture-related projects. One project is a market analysis expert system for oilseed commodities.

Another expert system provides diagnostics for dairy farms in countries where the dairy industry is largely undeveloped. Because veterinarians are often not readily available in many Third World countries, the expert system will help the people running the farms to diagnose problems with cattle.

Remote Sensing

The Canada Centre for Remote Sensing (CCRS) has

developed two hierarchical expert systems, the Analyst Advisor and the Map Image Congruency Evaluation (MICE) advisor.

"The analysis of Landsat images entails processing a very large amount of data," explains Dr. David Goodenough, chief research scientist at CCRS. "Our objective was to simplify this complex set of operations."

CCRS developed its own shell for this work and used a modular approach to building the expert systems.

"Instead of one giant system, which would be very difficult to use, we decided to have a number of systems working in cooperation," Goodenough says.

A resource manager can use Analyst Advisor to update forest inventory maps. Because the system analyzes satellite images for forest cover, the manager does not have to deal with the vision algorithms. Analyst Advisor will soon be in use in the British Columbia Department of Forestry.

"The advantage of having this expert system in the forestry office is that a lower level of skill is required to update the maps," Goodenough says. "In this way more information is disseminated at a lower cost."

Goodenough says the technology is readily transferable to other areas. He has had a graduate student working on a similar project for agricultural fields.

The Map Image Congruency Evaluation (MICE) system is used to compare maps and images. Because maps and images do not always agree, expert photointerpreters are required to make many judgement calls. MICE uses the same formal and heuristic rules to make decisions on reconciling the two data sources and on updating the map.

Several remote sensing projects are also underway at the University of British Columbia, some in conjunction with MacDonald Dettwiler and Associates. Dr. Alan Mackworth, a specialist in computational vision at UBC, says these projects deal primarily with the recognition of shapes on satellite images. Application areas include forest classification and identification of thermal upwellings in oceans.

Environmental Technology

MacDonald Dettwiler and Associates, the Canadian Atmospheric Environment Service (AES) and the Alberta Research Council have collaborated on a weather forecasting expert system. Called SWIFT (Severe Weather Intelligent Forecasting Terminal), the system will be used by AES forecasters.

The rules in the system embody the knowledge of an expert human forecaster. Without the aid of this advanced technology, it is virtually impossible for weather forecasters to identify, integrate and synthesize data in time to give more advanced warning of severe weather.

Dr. Marianne English, a scientist with the Alberta Research Council affiliated with the project, says the system's applicability will eventually be broadened to all kinds of weather. She says an expert system approach is ideal for handling meteorological problems which characteristically have "an awful lot of data and a solution which needs to be delivered in real-time."

As part of his Master's degree at the University of Alberta, Gordon McClymount turned his attention to tracking the flow of contaminants in groundwater. Expert ROKEY was the result — an expert system designed for users with a background in hydrogeology but who are not necessarily familiar with the complex computer models that describe groundwater flow.

Expert ROKEY (ROKEY is the name of a transport model) can be applied to a wide variety of contaminant sources, from sewage systems to gasoline spills. The system is in use at Environment Canada's Inland Waters Directorate in Burlington and Saskatoon, and at the Alberta Department of the Environment.

McClymount, who now works for Golder Associates in Calgary, is developing another expert system. Similar to Expert ROKEY, this expert system will help locate drilling waste sumps.

Drilling for oil produces fluid wastes which must either be stored on site or hauled off to another location. McClymount's system, based on a contaminant transport model, formalizes the process of sump location.

The expert system acts as the front end of the computer model — asking for data, evaluating results from the model, telling the user if the site is acceptable, and giving its reasons. Unlike Expert ROKEY, this system will not require the user to have an extensive background in hydrogeology.

Editors' note: As illustrated here, AI systems are currently being developed for the resource industry. The needs, however, are immense. But, as noted by Bill Havens in his editorial, the opportunities are there and challenging.

Research Directions for ICAI in Canada

by Philippe Duchastel

Directions de recherche en EIAO au Canada

RÉSUMÉ: En plus de donner une idée des différentes activités de recherche en enseignement intelligemment assisté par ordinateur (EIAO) en cours au Canada, cet article discute plus généralement de la nature de la recherche en EIAO et examine vers quelle méthodologie de recherche et quels problèmes une attention particulière pourrait être portée en vue de promouvoir cette recherche dans un contexte national.

Introduction

In [6], I argue that current instructional design technology is inappropriate for the evolving needs of the educational and training communities, particularly in view of the possibilities of intelligent computer-assisted instruction (ICAI) for deeply engaging the student. Current efforts in ICAI are strongly influencing our view of learning and teaching. This influence will be felt strongly in our approach to instructional design.

Philippe Duchastel is the director of Laboratoire d'Intelligence Artificielle en Education at Université Laval.

Instructional systems design (ISD) will gradually be replaced by a modified paradigm centered on learning opportunities and characterized by learning environment design (LED). This new view of design might well have profound implications on how we consider educational technology and how we engage in the process of educational design.

This view of design will underlie the educational and training systems of the 1990's and therefore it is important for us to consider what prospective R&D efforts currently seem most promising within this context.

While [6] was focused on instructional design *per se*, the arguments apply with the same force to the hard technologies based on that design approach. They apply in particular to computer-assisted instruction (CAI) and computer-based training (CBT), technologies with which ICAI is often contrasted in order to bring out its own particular approach and its own potential [1,4].

This contrast can perhaps be illustrated by briefly considering SOPHIE, an early ICAI system which took the form of a computer-simulated electronics laboratory, the goal of which was to provide the student with practice in troubleshooting faulty electronic equipment [2]. A traditional CBT approach requires a complete specification of the anticipated moves and queries of the student, a task which is never completely successful and which therefore leads to a rather directive teaching strategy. The student is lead through reasoning processes which may or may not correspond to his particular needs and interests at that moment. SOPHIE, on the other hand, provided a reactive learning environment which encouraged the student to explore new ideas and formulate hypotheses, thus allowing the student to learn in his own way how to solve the problem. The system offered an opportunity "for providing the student with detailed logical analyses of correctness of his hypotheses just when the student is most likely to be interested in such feedback" [1]. An inquiry approach to learning, in which the student directs the learning interaction, was at the heart of the design philosophy of SOPHIE, and is likely so for the greater part of ICAI as well.

This contrast illustrates the potential power of ICAI and of the LED approach generally, especially in terms of process and contextual outcomes.

Research Directions

Given that we want to bring about such a technology, which particular directions do we pursue? Two fundamental aspects are involved. First, what type of research is needed? What research methodology is most appropriate in this area? Second, which specific research questions are important? What focus should the research have?

With respect to the type of research, it can still be argued that the royal road to progress in AI, and thus in the subfield of ICAI as well, lies first and foremost in the design of prototype systems. It is in the development of such prototypes that the practical requirements and constraints of computational intelligence are felt. It is especially here that limitations of current methods and tactics are perceived, thus leading the way to theoretical insights from which novel strategies emerge. AI is first an engineering science, one which builds more on practical experience gained through trial and error than on theoretical analyses derived from the cognitive sciences, the object of which it is trying to model.

This is not to imply that the cognitive sciences play only a minor role in the progress of AI. Indeed, the methodology which is allied to prototype design in AI, the one that conceptually underlies it, is cognitive process modeling. Any AI system is a computer model of some human process, and a good understanding of that process is therefore required if the AI system is to prove successful. Cognitive modeling is explicitly illustrated in the field of ICAI by the pioneering work of Collins' research group at BBN in the 1970's [3,5,10]. An explicit attempt at analyzing the Socratic tutoring processes engaged in by human tutors was the basis of this work aimed at eventually incorporating these processes in a computer tutor.

Thus, with respect to methodology, two aspects are seen as crucial for progress within ICAI. First, an emphasis on prototype development is essential; actual work on the development of intelligent tutoring systems is needed. Second, development work of this kind will profit greatly from explicitly modeling the cognitive processes in situations such as those being mimicked. It is easily seen that such research can profitably involve interdisciplinary teams of researchers with specific backgrounds in either computer science or the cognitive sciences, especially psychology.

Turning now to the focus of the research, one sees an initial difficulty in that the field of AI is very large and varied (ranging from robotics, to expert systems, to natural language understanding systems). Furthermore, ICAI systems developed to date show a great variability in their structure and their functioning [9,11], as well as in their particular goals.

Even if ICAI activity is very diverse, an intelligent tutoring system (ITS) is generally thought to be comprised of four complex and interrelated components: domain expertise, student model, tutorial heuristics, and interface. The tutorial component must reach its decisions in a finely understood context, the basis of which is the dynamic student model evolving as the learning session progresses. This model in turn is a differential model which constantly contrasts what the student knows with what is considered expert knowledge or skill in the domain. Finally, language and communication processes are largely knowledge-based and generally cannot be treated outside of a semantic context. Thus, an ITS attains a level of performance as a whole, even if specific instances of ITS's concentrate their research attention on particular issues within the ICAI field.

Another concern is the specificity of the research foci within the larger context of AI. Some research concerns are found at a broad level in AI. Knowledge representation and language processing are two of these. On the other hand, tutorial processes and student modeling are somewhat more specific to ICAI. This is not totally the case, however, as user modeling is becoming more and more a concern in such areas as information retrieval and dialogue systems. While ICAI has its own particular concerns and its own context, it seems that more sharing is occuring with other application areas of AI generally. This bodes well for ICAI.

Canadian ICAI

It is useful here to note some of the more active research projects being conducted in Canada. At Simon Fraser University in British Columbia, a well-funded project named DOCENT is developing a knowledge base

on teaching which will serve as the basis for an expert system to help teachers plan their classroom lessons. In Calgary, the Alberta Research Council and the company Computer-Based Training Systems is developing an expert computer-managed learning system which will be initially tested within five Canadian educational institutions.

In Victoria, the company Softwords is incorporating production rule formalisms to enhance the delivery of instruction within traditional CAI training systems. Likewise, the BIOMEC project at Laval University is incorporating a production rule approach to instruction within an interactive videodisk-based learning environment.

One of the more mature Canadian projects is underway at the ARIES Laboratory at the University of Saskatchewan in central Canada. This project is developing a large intelligent tutor named SCENT to support the learning of Lisp. The design of a number of smaller tutors, either game-based (in geography and nutrition) or simulation-based (in optics) is also underway at the LIAE Laboratory at Laval University in Quebec.

In Montreal, a private firm has developed the SCARABEE natural language software to help students write adventure-based stories. In Ontario, a university and software firm consortium is developing a similar natural language-based writing package. At the University of Montreal, the HERON project aims to develop, among other things, an intelligent advisor for a word processing package (see *Canadian Artificial Intelligence*, April 1988). It is this group at the University of Montreal that hosted the International Conference on Intelligent Tutoring Systems in early June.

Finally, an international project involving the National Research Council of Canada in Ottawa, the University of Leeds in Britain, and two groups in Quebec (a Price Waterhouse software affiliate and Laval University) aims to develop an intelligent advisor system to assist data modeling in database design.

These brief notes on some of the major Canadian activities in ICAI provide an indication of the variety of work being conducted in this area in the country. They do not do justice, however, to the complete range of activities pursued in ICAI in Canada, for there are a number of other projects of varying size and depth which have not been mentioned here. A more comprehensive listing has recently been compiled by Jim Greer of the ARIES Laboratory for the Associate Committee on Instructional Technology of the National Research Council of Canada.

Context for Prospective ICAI Research

It was argued above that the principal way to progress in ICAI research is to undertake the development of ICAI systems. It is in designing concrete prototypes that the various research issues alluded to above are met head on and handled in some particular manner. It is by considering the eventual limitations of the prototype that more specific research issues are identified; these in turn serve to specialize the research agenda. What is aimed for in the long term is an ICAI design technology similar in scope to that of CAI or to the more recent one concerning expert systems.

One may wonder whether there are particular choices which should be made in getting such a research effort underway. Two issues arise. The first relates to the choice of domain and of system type for initial efforts. The second concern relates to relations with the research community currently involved in ICAI research.

ICAI systems have been and are being developed in a variety of domains, including geography, electronics, medicine, mathematics, geometry, physics, and computer programming. They have involved different types of expertise, such as factual knowledge, the application of principles, particular cognitive skills, and reasoning strategies. They have involved tutorial dialogues, expert systems, instructional games, and advisory (active help) systems. They have been applied to formal educational tasks, to technical training, to on-the-job training, and to informal learning environments. Thus the existing prototypes cover a wide range of areas and styles. This diversity is not unlike that found in traditional CAI, where different styles of application have emerged over the years.

Whether one set of application areas or one set of styles can be considered more important than others in social terms is a very difficult question. None of them individually have received the degree of research attention which would lead us to believe that the area is well understood. Each could profit from much more extensive development efforts.

In terms of importance, the choice of an area is practically a matter of personal taste. Of course, in terms of contextual justification, considerations such as funding opportunity or institutional mission may be present. These, however, do not relate directly to the research issues themselves, which are context independent.

If we agree that the future will see a trend involving less formal education and training and more spontaneous learning of an informal nature, a certain concentration on advisory systems may be worthwhile rather than on formal ITSs. Examples of such systems are the Canadian ADVISOR project headed up by the National Research Council in Ottawa or the NLS-SCHOLAR system for a complex computer editor [7]. The line between the training and other components of such AI systems as expert systems, advisory systems, and simulation systems is often rather blurred. Furthermore, the very affinity of ICAI with inquiry learning and informal learning environments gives a very open character to even dialogue-based tutorial systems. A choice of style, therefore, does in the end remain rather arbitrary.

The second contextual concern involves the research community currently active in ICAI. Any prospective research effort must attempt to tie into the existing net of on-going efforts in order to profit from the experience already generated. Most efforts to date have been based in the U.S., although Great Britain also has a certain tradition in the area. Particular centers of interest are Bolt, Beranek & Newman in the Boston area, the University of Massachusetts at Amherst, Carnegie-Mellon University and the University of Pittsburgh, all in the eastern U.S., and Xerox Parc and Stanford University on the west coast (The Institute for Research on Learning, in Palo Alto is a recent addition). Other countries such as France, Germany, and Italy also have groups working in the area.

Conclusion

AI is a rapidly growing field of research, for a number of countries are investing substantial research funds in what is perceived to be one of the most important technological areas of the future. AI technology will have

an impact on most sectors of human activity, including education and training. ICAI seems to be poised to revolutionize educational technology and offers the promise of providing very sophisticated learning systems which may well greatly increase both the effectiveness and the pleasure of learning. The research which is called for today in the field of ICAI will set the framework for this new technology. Exactly what shape this technology will take is difficult to predict. Even what particular priorities should be taken up is far from clear. What does seem certain, however, is that experience and expertise is not gained by staying on the sidelines and merely observing the results of others. The technology is there waiting to be pushed forward, something that is starting to happen in Canada because of the determination of individual research groups and software firms convinced of the potential they see in ICAI.

References

[1] Brown, J.S., Burton, R.R. and Bell, A.G., "SOPHIE: A step toward creating a reactive learning environment" *International Journal of Man-Machine Studies, 7,* 1975, pp. 675-696.

[2] Brown, J.S., Burton, R.R. and de Kleer, J., "Pedagogical, natural language and knowledge engineering techniques in SOPHIE I, II and III" *Intelligent Tutoring Systems* Sleeman, D. and Brown, J.S.(editors). Academic Press, New York, 1982, pp. 227-282.

[3] Burton, R., "Diagnosing bugs in a simple procedural skill" *Intelligent Tutoring Systems* Sleeman, D. and Brown, J.S.(editors). Academic Press, New York, 1982, pp. 157-183.

[4] Carbonell, J., "AI in CAI: An artificial intelligence approach to computer assisted instruction" *IEEE Transactions on Man-Machine Systems, MMS-11:4,* 1970, pp.190-202.

[5] Collins, A., "Processes in acquiring knowledge" *Schooling and the Acquisition of Knowledge*, Anderson, R.C., Spiro, R. and Montague, W. (editors). Lawrence Erlbaum Assoc., Hillsdale, New Jersey, 1977, pp. 339-364.

[6] Duchastel, P., "Designing intelligent learning environments" To appear in *Designing New Systems and Technologies for Learning*, Mathias, H., Budgett, R. and Rushby, N. (editors). Kogan Page, London, U.K.

[7] Grignetti, M.C., Hausmann, C. and Gould, L., "An 'intelligent' on-line assistant and tutor — NLS-SCHOLAR" *AFIPS conference proceedings*, 44, 1975, pp. 775-781.

[8] Kintsch, W., Miller, J. and Polson, P., *Method and Tactics in Cognitive Science*, Lawrence Erlbaum Assoc., Hillsdale, New Jersey, 1984.

[9] Sleeman, D. and Brown, J.S., (editors) *Intelligent Tutoring Systems,* Academic Press, New York, 1982.

[10] Stevens, A. and Collins A., Goldin, S.E., "Misconceptions in students' understanding" *Intelligent Tutoring Systems*, Academic Press, New York, 1982, pp. 13-24.

[11] Wenger, E., *Artificial Intelligence and Tutoring Systems*, Morgan and Kaufmann, Los Altos, California, 1987.

AI Research and Development at Applied AI Systems, Inc.

by Takashi Gomi

Recherche et développement chez Applied AI Systems, Inc.

RÉSUMÉ: Applied AI Systems, Inc. (AAI) est une firme de recherche en IA appliquée qui s'occupe aussi de développement de systèmes. Depuis 1983 AAI a terminé des aperçus de recherche, des études de practicabilité, et le développement de logiciels dans plusieurs sous-domaines de l'IA pour des organisations gouvernementales, industrielles et étrangères. Cet article présente les projets-clés entrepris chez AAI.

Applied AI Systems, Inc. (AAI) is an applied AI research firm engaged in both applicative AI research and system development. AAI's staff provides expertise in the research, design, implementation, and deployment of practical knowledge-based systems in the fields of government operations, engineering, transportation, communication, and business. AAI organizes and focuses the know-how obtained from world-wide expertise (approximately 1,700 personal AI contacts established throughout the world to date), to prepare practical problem-solving methods for its clients.

Since its establishment in January 1983, AAI has completed surveys, feasibility studies, and application software development in several subfields of AI for various governmental, industrial, and foreign organizations. During this five year period, the company has worked with over sixty Canadian and international clients, including Transport Canada, The Ontario Ministry of Transportation and Communications, Nippon Telegraph and Telephone (NTT), The Canadian Department of External Affairs, The Canadian Department of Communication, The Canadian International Development Agency (CIDA), Revenue Canada, The Canadian Department of Energy, Mines and Resources, Arctec Canada, The Japanese Institute for Future Technology, and Advanced Telecommunications Research International (Japan). Projects carried out have included a study of the applicability of expert systems technology to the Canadian Transportation sector, a feasibility study on AI technology in marine operation, a survey on the application of voice interactive technology, the development of a knowledge-based scheduling system for a municipal transit operation, the development of an expert system for highway transportation, a survey of advanced AI R&D centres throughout the world, and research on the frontier of advanced natural language understanding. In addition, a number of knowledge-based system developments and research projects are currently in progress.

AAI also introduces affordable American and European AI products to the Canadian market, acting as agents for the most practical and technically innovative products encountered in the course of its global research. (See advertisement on the back cover of this magazine for AAI products.)

Takashi Gomi is the president of Applied AI. He has a formal background in electrical engineering, and over 25 years experience in computer science and AI.

Technical Advisors

AAI has established contacts with a large number of researchers, project administrators, government officials, engineers, manufacturers and users involved in AI R&D in Canada, U. S., Britain, Continental Europe, Japan and other Asian countries. Since it emphasizes the quality of information it deals with, the company routinely exchanges information with these contacts. In addition, over the past several years, AAI has obtained the services of the following world-class individuals in the capacity of technical advisors, lecturers, consultants, or as members of project teams as indicated:

- Dr. Thomas Adams, Advanced Decision Systems, California: Space-based AI Systems and Data Fusion (consultant, project member)
- Professor Kenneth Bowen, University of Syracuse: Logic Programming (lecturer)
- Professor Michael Brady, University of Oxford: Autonomous Systems and Real-time AI Systems (advisor, consultant)
- Professor Rodney Brooks, Massachusetts Institute of Technology: Intelligent Robotics and Autonomy (consultant)
- Professor Paul Cohen, University of Massachusetts at Amherst: Knowledge Representation (consultant)
- Professor Tim Finin, University of Pennsylvania: Intelligent Interfaces (lecturer)
- Dr. Lucca Gilardoni, Quinary Systems, Milano: Qualitative Reasoning in Industrial Systems (consultant)
- Dr. Peter Hammond, Imperial College of Science and Technology, London: Prolog and Logic Programming (lecturer)
- Professor Robert Kowalski, Imperial College of Science and Technology, London: Logic Programming, Co-inventor of Prolog (advisor, lecturer)
- Professor Victor Lesser, University of Massachusetts at Amherst: Distributed and Real-time Artificial Intelligence and Real-time AI Systems (advisor, lecturer, consultant)
- Professor Drew McDermott, Yale University: Logic (advisor, consultant)
- Professor Alan Mackworth, University of British

Columbia: Computational Vision (lecturer)
• Professor Hiroyoshi Obara, Waseda University: Natural Language Processing (advisor)
• Professor Franz Oppacher, Carleton University: Schankean Natural Language Understanding and Expert Systems (lecturer, advisor, project member)
• Professor Alan Robinson, University of Syracuse: Logic and Logic Programming (lecturer)
• Dr. Stan Rosenschein, SRI International: Autonomy and Logic (consultant)
• Professor Elliot Soloway, Yale University: Intelligent Tutoring Systems (lecturer)
• Dr. Paul Tuan, SRI International: Computerized Traffic Management Systems (project member, advisor)

Projects

The following is a list of key projects undertaken by Applied AI Systems, Inc. (AAI) in the field of applied artificial intelligence.

Expert Systems in Canadian Transportation

In 1985, in conjunction with Cognos Incorporated, AAI conducted an investigation into how expert systems technology could be applied to the Canadian transportation sector. The study covered the following topics:

• principles of expert systems
• how they can fit into the transportation sector
• what they can and cannot do
• the cost of expert system development
• example transportation- oriented expert systems
• details of the process of applying expert systems to transportation
• market trends in expert systems and related technologies
• available tools for implementing expert systems
• Canadian expertise in expert systems and the supply of knowledge engineers.

A detailed analysis of potential applications resulted in the identification of approximately 30 likely candidates. Several of the applications which scored high in this analysis have since become projects, attesting to the accuracy of the analysis. Over 1,400 copies of the report have so far been requested by both government departments and industrial organizations in Canada and at least seven other countries.

Contract Evaluation Expert System

A contract evaluation expert system is being developed for use by officials in Canadian government departments and agencies. The project was submitted in 1987 as an Unsolicited Proposal, and is currently being developed with the cooperation of the Canadian International Development Agency (CIDA).

Knowledge of the regulations governing the awarding of contracts is very experience-intensive, and is influenced by many different factors. Different actions are required depending upon the circumstances surrounding the contract (e.g. the amount of the contract, whether ex-government officials are involved, the duration of the contract, etc.). The principal aim is to maintain consistency in the awarding process, regardless of the government officials involved in the assessment. An expert system is an ideal tool with which to achieve this. An extensive report generator documents the entire assessment process, as well as the results of the assessment. In the CIDA context, an expert system would assist less experienced contracting officials in different areas of the world to apply a uniform set of rules when awarding contracts. The system will thus contribute to the decentralization of CIDA's operations. Such decentralization is a general move in many government agencies.

Voice Interactive Technology

One of the important extensions of AI systems in the real world will be the addition of speech and natural language (NL) capabilities. Commissioned by Transport Canada in 1987, this study conducted a survey of voice (speech) and NL technologies as they were related to transportation applications. An overview of the current state-of-the-art in voice technology was contributed by Dr. Melvyn Hunt of the National Research Council of Canada. This was accompanied by an extensive survey of practical NL research in academia and industry.

A series of interviews was conducted with Transport Canada officials involved in the application of voice interactive technology to identify current application cases and to investigate further potential for these technologies in the transportation sector. Available voice and NL products were also surveyed and listed, in order to assist those planning to start a project using these technologies.

AI in Marine Operation

A study commissioned by Transport Canada in 1986, jointly with Cognos Incorporated, examined opportunities for applying expert systems and more advanced AI techniques in the operation of ocean-going ferries and other similar ships. A detailed case study was conducted on CN Marine's M. V. Caribou for the automation of navigation, fuel-efficient operation, and on-board in-operation maintenance. Application of real-time knowledge-based system (RTKBS) formalisms was also outlined for a number of on-board automation applications.

As to the navigation application, a series of test runs by the actual crew of M. V. Caribou (3 captains and 3 first officers) using the U. S. Coast Guard's simulator with a large scale graphic front-end was used to acquire knowledge for port entry. Test runs were conducted on the scenario of M. V. Caribou, one of the largest ocean going ferries known, attempting to enter Port-aux-Basques (one of the narrowest, most treacherous and climatically demanding ports in Canada) under various environmental (wind, current, fog) and on-board (load, engine conditions) situations. Excellent results were obtained from these knowledge engineering sessions. A detailed analysis of the process resulted in a set of suggestions for what must be done to eventually automate the complex port entry operation using various AI techniques (some 40 identified), ranging from sensor fusion to qualitative reasoning. The results from this study were published as a Transport Canada report.

AI Research in Japan, Britain, and Europe

From the fall of 1985 to the spring of 1987, under the Technology Inflow Program of External Affairs Canada, AAI visited a number of applicative AI R&D centres in Japan, Britain, and Europe. These visits resulted in a large number of technical and business contacts in those parts of the world. The results of the visits were

summarized in three reports, which have been widely distributed in Canada.

Foundation for Advanced AI Applications

In 1988, AAI received a research contract from a major telecommunications company in Japan on the development of an advanced AI system concept. For competitive reasons, the details of the project cannot be revealed. However, the research involves consultation with the world's top AI researchers, in both academia and industry, in the United States, Britain, and Europe. The contract includes the delivery of a report including a summary on the state of the art in the field and a plan for further development.

Sixth Generation Computing

Since early in 1987, AAI has been conducting in-house research on sixth generation (6G) computing, loosely defined as AI beyond knowledge processing, or non-symbolic AI. The first phase of this study was a survey of related research activities throughout the world today. Topics investigated included the Japanese Human Frontier Science Program (HFSP) and the Advanced Telecommunications Research (ATR) project as its first implementation, and the European Commission's BRAIN project. Topics also included the emerging neural networks or neural computer concept actively pursued in the United States, Europe, and Japan. This first phase has recently been completed, and the results presented at the 2nd International Symposium on Artificial Intelligence and Expert Systems held in Berlin in the summer of 1988. For the next phase, AAI will conduct experiments using neural network (NN) models.

Advanced Natural Language Research

In late 1987, AAI received a research contract from the Institute for Future Technology (IFTECH), a dominant think-tank in Tokyo sponsored by the Japanese Agency for Science and Technology. IFTECH's mandate is to research and report on trends, and to recommend R&D plans to the Japanese government and industry for dealing with long-term science and technology issues. They run approximately 70 research projects annually, in various fields.

The project AAI received was to review the natural language (NL) research, which is today dominated by research conducted from a linguistics perspective, beyond the current state of the art. AAI has since conducted a detailed survey of practical NL research throughout North America. The study includes a detailed account of several schools of Schankean NL research and multi-model intelligent interface research in the United States. A substantial report was delivered to the satisfaction of the client.

Fusion of Symbolic and Neural AI Paradigms

Advanced Telecommunications Research International (ATR) is a research organization established jointly by the Japanese Ministry of International Trade and Industry (MITI) and Ministry of Post and Telecommunications (MPT). It is also one of the first research centers in the world for sixth generation computing, implementing many of the concepts outlined in the Japanese Human Frontier Science Program initiative announced to the world in 1986. ATR's current approved budget to 1994 is about Cdn$750 million equivalent. Greater funding is expected for the second 7.5 year phase that follows.

AAI has received a study contract in 1988 from one of the four research groups at ATR, the Interpretive Telephony Research Laboratories. Their goal is to literally build an AI-based translating telephone system between Japan and other countries, starting with the United States. This necessitates accelerated R&D in speech technology, NL, and machine translation (MT). One of the current dominant research themes at the Laboratories is the application of neural network (NN) paradigms to NL Understanding. AAI's first research assignment is to propose various approaches for interfacing, or the fusion of, symbolic and neural AI paradigms. The second role is to inform the group of significant research activities in NN-based NL methods throughout the world.

AI System Applications within Governments

Since 1983, AAI has won several AI R&D projects from federal and provincial government agencies in Canada. The experience of working on these projects for government agencies resulted in a series of survey papers on the use of AI in governments throughout the world. The research results have been reported in presentations given in Europe (Germany), Washington (at the Third Expert Systems in Government Conference, ESIG-87), and Ottawa. It is an open-ended survey project, and AAI intends to continue it into the future.

Knowledge-Based Vehicle Dispatcher's Assistant

As a result of a proposal submitted to Supply and Services Canada under the Unsolicited Proposals program, a contract was awarded in 1987 for the development of a prototype expert decision aid to support truck dispatchers. The project is funded and managed by the Transportation Development Centre of Transport Canada.

The prototype which is currently under development is intended to support truck dispatchers by capturing the skills and knowledge used by an expert in the actual dispatching of the trucks, and thus frees the expert's time for more complex and creative tasks. Canadian Liquid Air (CLA), Hamilton, Ontario has agreed to participate in the project as the test site.

The development is on a 386 computer using the NEXPERT shell.

Spacecraft Autonomy Management

In 1984, AAI was granted a prototyping contract from the Canadian Department of Communications for an on-board autonomy management system. This followed one and a half years of thorough study by AAI of on-board fault tolerance requirements and methodologies, and autonomy application requirements; a survey of on-board AI system techniques; and, functional design of an on-board autonomy system based on a hierarchically organized, multiple expert system architecture.

AAI developed a computational framework for the on-board intelligent system. The framework combined an evidential reasoning model based on the Dempster-Shaefer data fusion paradigm and a goal driven reasoning mechanism. This combination is starting to appear today in a number of real-time AI systems for aerospace and industrial applications. The model was implemented in some 3,000 lines of MPROLOG code on a VAX-11/780. A series of knowledge acquisition sessions with a communication satellite design expert and satellite operations specialist was conducted to elicit knowledge of

recovering a satellite in serious anomaly.

A series of experiments using the prototype and simulated on-board operational environment was conducted. It has shown that a serious error in a satellite attitude control mechanism caused by a set of on-board anomalies can be diagnosed and successfully recovered from in about 13 minutes, compared with the several hours required by a team of human experts. AAI has since been approached by Sperry's Space Division to develop identical experiments for them.

Expert System Demonstration in Urban Transit

A contract was awarded as a result of a competition in 1986 by the Ontario Ministry of Transportation and Communications (MTC) to develop a prototype knowledge-based system to demonstrate the applicability of the technology in transit operation. For the execution of this project, AAI joined forces with Cognos Incorporated and with SRI International of Menlo Park, California. A prototype was developed on an IBM PC/AT using Arity Prolog. The Kitchener Transit Centre was selected as the demonstration site and spare-board management as the demonstration task. Spare-board management deals with solving emergencies in transit dispatching (e.g. loss of a driver for a scheduled run by reason of sickness). The prototype was designed entirely around the existing set of paper documents in order to make it truly useful in the field. This decision resulted in an expert system architecture which is totally different from conventional interactive models.

After the completion of the demonstration project, AAI has undertaken the task of developing the prototype further to make it a deployable system. Considerable effort has been expended to improve user and database interfaces, and to make the system robust, carefree and fast. This phase of the development has been reported at the International Workshop on Artificial Intelligence for Industrial Applications held in Japan in May 1988.

Artificial Intelligence Work at MacDonald Dettwiler and Associates Ltd.

by Max Krause

Travail en intelligence artificielle chez MacDonald Dettwiler and Associates Ltd.

RÉSUMÉ: Au cours des dernières années, les applications de l'IA sont devenues une composante de plus en plus importante du programme de recherche de MacDonald Dettwiler. Deux des projets en cours, PREFORMA et SWIFT, touchent à l'un des domaines traditionnels de la compagnie, la météorologie. Un troisième projet, FEX, s'intéresse à la cartographie à partir de l'espace. Le tout dernier projet, STAR, utilise l'expertise de MacDonald Dettwiler dans un nouveau domaine, celui de l'entretion de systèmes électro-mécaniques. L'on s'attend à ce que ce domaine pour l'instant terrestre devienne utile pour la recherche spatiale.

Introduction

MacDonald Dettwiler and Associates Ltd. (MacDonald Dettwiler) is a world leader in the technologies of remote sensing, particularly in the reception and processing of imagery from space. MacDonald Dettwiler has developed and installed systems world-wide for the reception and processing of data from earth resource satellites such as Landsat and SPOT, from meteorological satellites in both geosynchronous and near-polar orbits, and from synthetic aperture radar (SAR) satellites. Since its founding in 1967, MacDonald Dettwiler has striven to anticipate, to develop and to use state-of-the-art technology in its products. AI is one of those technologies which appears to be poised to become an essential part of any leading-edge company's tool kit, as it has done at MacDonald Dettwiler.

Applications of AI have become an increasingly important part of MacDonald Dettwiler's research program over the past few years. With a current complement of five AI workers, the knowledge based systems area comprises a significant and growing fraction of our research staff and research effort.

Two of the projects undertaken to date, PREFORMA and SWIFT, have been in one of the company's traditional application areas, meteorology. A third project, FEX, is applied to the problem of mapping from space, which is becoming an increasingly important area of business for MacDonald Dettwiler. The latest project, STAR, applies MacDonald Dettwiler's expertise to a new domain, maintenance of electro-mechanical systems. Although the current application domain of STAR is terrestrial, we anticipate the need for this type of expertise in space as well as in industries on earth. As a member of Canada's Space Station Project team, MacDonald Dettwiler is particularly concerned with fostering and enhancing Canadian skills in AI for potential future application in that project.

The following short accounts summarize some of the AI work undertaken in the recent past at MacDonald Dettwiler.

Max Krause is an AI researcher employed by MacDonald Dettwiler and Associates. He holds a B.Sc. in Psychology and an M.Sc. in Computer Science.

PREFORMA

The PREFORMA (PREcipitation FORecasting Meteorologist's Assistant) project was undertaken to test the feasibility of achieving expert level performance by a knowledge based system in the field of site-specific short term weather prediction. The approach taken was to model the conceptual categories and the inference methods applied to this problem by human weather forecasters. This approach was motivated by our perception that this problem was being solved with reasonable success by human forecasters in day to day operations, and that a numerical approach to the problem demanded resources that were beyond the means of most potential users of such a system. In order to limit the

PICON®

delivers the promises of Artificial Intelligence and Expert Systems, Now.

PICON goes beyond expert system shells by providing a solution to the problem of large scale coordinated control.

When interfaced to a distributed control system via RTIME, PICON allows unprecedented productivity where it counts: in applications to real world problems, on line and in real time.

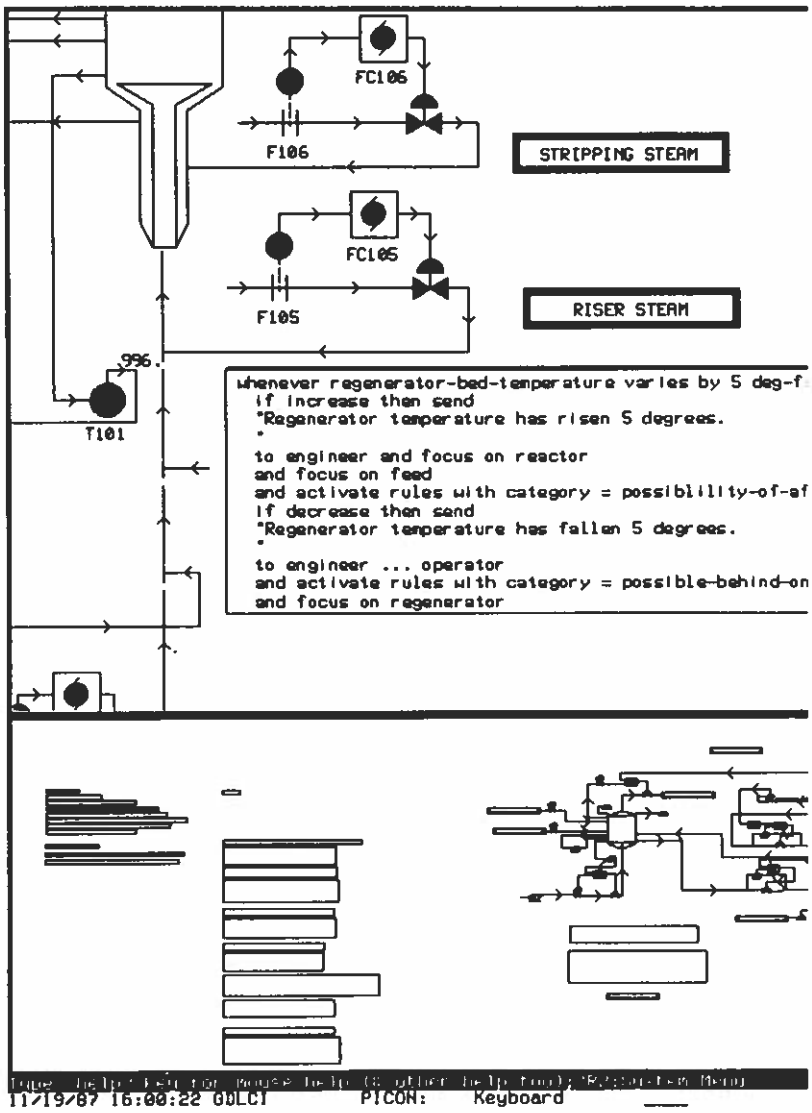

Process Intelligent Control (PICON) *is application system. It is developed from and process control.*

PICON is distributed and supported by LISP CANADA Inc., a member of the GigaMos Group companies, LISP CANADA also distributes and supports the LAMBDA computer and other products

GigaMos Systems Inc.

a leader in the application of artificial intelligence.

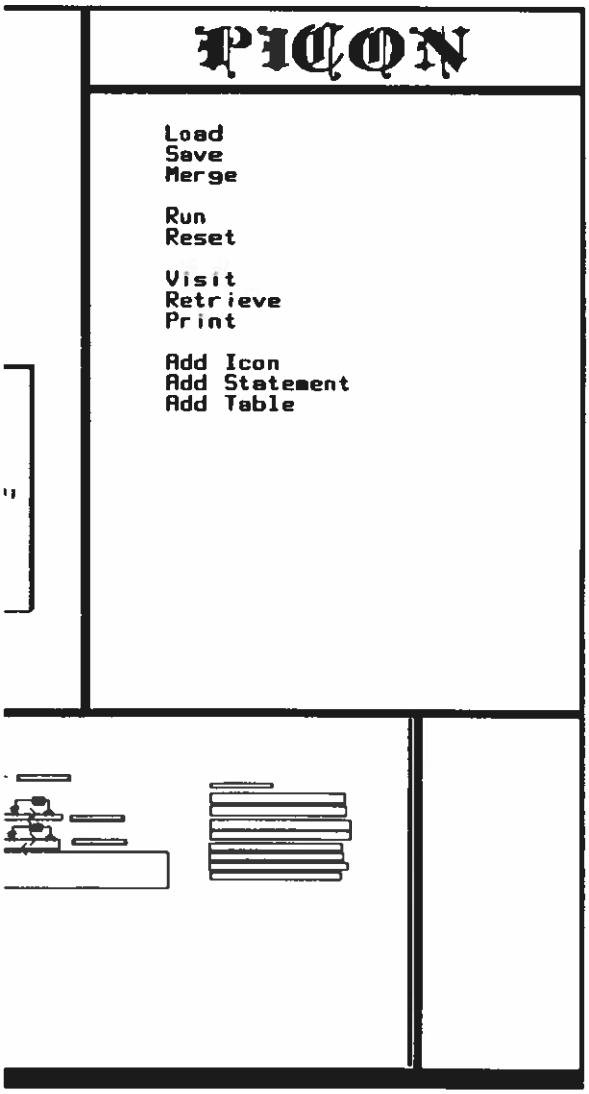

ule-driven, definition-based
erience in artificial intelligence

Here are some of the reasons why PICON is intelligent and productive:

- NATURAL LANGUAGE DESCRIPTION OF THE EMPIRICAL AND ALGORITHMIC UNDERSTANDING OF THE PROCESS to allow the user to be productive without programming.
- HEURISTIC SOLUTION TO COMPLEX PROBLEMS to deal with large amounts of data and complex, time-sensitive decision.
- LIVE PROCESS GRAPHICS WITH VALUES ENTERED FROM THE CONSOLE to test the behavior of the rule base before using real data from the process.
- EASY USE OF GRAPHICS to allow the user to represent the process as he is used to seeing it.
- HIGH RESOLUTION GRAPHIC CONSOLE WITH MOUSE AND KEYBOARD providing state of the art user interface.
- WINDOWS AND HELP FEATURES to make the interface user friendly.
- SYSTEMATIC AND SELECTIVE DOCUMENTATION to allow disciplined implementation.
- MENU-DRIVEN, GUIDED; RULE-BASE, DATABASE ENTRY to minimize the effort needed in application.
- A VARIETY OF COMPUTERS AS DELIVERY VEHICLES to allow the best systems solution.
- INDUSTRY STANDARD INTERFACES TO DISTRIBUTED SYSTEMS to ensure rapid systems integration.

LISP Canada Inc.

1000 Roche Boulevard
Vaudreuil, Québec J7V 6B3

Telephone: (514) 455-5605
FAX: (514) 455-0400

complexity of the problem, we took the advice of the Atmospheric Environment Service and restricted it to 12-hour prediction of the occurrence, phase, timing and intensity of precipitation for a single site.

After about two person-years of effort, PREFORMA has a very good user interface which supports the analysis of weather maps in a way that would speed up the analytical work that an operational forecaster performs. The meteorological reasoning component of the system can perform the automatic projection of future positions of cold and warm fronts using rules of thumb employed by forecasters, although with lower accuracy than human forecasters. This discrepancy in performance is due to the difficulty of performing on the computer the complex visual pattern recognition tasks that human forecasters use to project the future positions of weather features.

Participants: Guy Druce, Max Krause, Peter George

SWIFT

The SWIFT (Severe Weather Intelligent Forecast Terminal) project is a joint effort by MacDonald Dettwiler and the Alberta Research Council for the Atmospheric Environment Service of Environment Canada. SWIFT is a prototype system designed to provide a forecaster with a workstation which aids in the tasks of organizing and interpreting the large amounts of information relevant to making same-day predictions of severe weather in the Canadian prairie provinces. The system makes use of both a statistical index of potential convective activity (the Sc4 Index developed by Geoff Strong of the Alberta Research Council) and knowledge of observed weather features to produce predictions of the location and direction of motion of areas where severe convective weather is likely to occur within the current day.

The "intelligence" of SWIFT lies in several areas:

- a module which parses the remarks section of SA reports to extract information about weather events such as cloud and precipitation
- a module which uses knowledge of observed weather features and events to refine the interpretation of a statistical predictor of storm initiation in the eastern foothills of the Rocky Mountains (the Sc4 index)
- a sub-system which captures an internal (frame-based) model of a weather chart analysis of features and fields — the results of subjective or objective analyses are thus available to other sub-systems which can use such knowledge to reason about future developments
- spatial (geographic) and temporal reasoning ability.

SWIFT is implemented in Interlisp-D, LOOPS, and OPS4 on a Xerox 1186 Lisp computer. The reception and formatting of incoming data and the calculation of the Sc4 index is done on a host computer (currently a DEC VAX); the workstation software accesses the host disk for data as needed.

Participants: Max Krause, Peter George (MacDonald Dettwiler), Doug Konkin, Ken Gamble, Julia Driver (Alberta Research Council)

FEX

FEX (Feature EXtraction) is a knowledge-based system for automatic extraction of planimetric features from remote sensing imagery. Topographical planimetric features include natural surfaces (rivers, lakes) and man-made surfaces (roads, railways, bridges). In conventional planimetric feature extraction, a photointerpreter manually interprets and extracts features from imagery on a stereoplotter. Visual planimetric feature extraction is a very labour intensive operation. The advantages of automating feature extraction include:

- savings in time and labour
- accuracy improvements, and
- planimetric data consistency.

Present research efforts at MacDonald Dettwiler have concentrated on the extraction of road-like features in SPOT imagery. Numeric techniques by themselves would not achieve the consistency and accuracy of human photointerpretation. In order to detect and locate road features, FEX incorporates heuristic knowledge of photointerpretation. Some of the knowledge associated with roads is as follows:

- Spectral — Road segments in the imagery have higher spectral contrast than the immediate surrounding neighbourhood. Road segments tend to have a local uniform spectral distribution.
- Geometric — Parallel line structures give hints on the location of potential road segments. Locating parallel lines from the detected edges in the imagery requires the extraction of a hierarchical network of abstract geometric structures.
- Structural — Roads form networks and intersect other road segments.
- Contextual — Roads are not located in water. Houses and buildings are usually located near roads. A bridge is a road segment that crosses over a water body.

FEX combines techniques from image processing, remote sensing and artificial intelligence for automatic feature extraction. The FEX prototype has incorporated some of the spectral and geometric knowledge. This prototype has been able to reduce commission errors and omission errors to less than 30%. Research indicates that the error levels can be reduced even further by including structural and contextual knowledge.

Numerical processing in FEX is handled by VAX Pascal, while symbolic data manipulation is managed by Quintus Prolog. FEX is implemented on a DEC microVAX.

Participants: John Zelek, Ben Yee

STAR

The STAR (STrategic Advanced Research) project is a cooperative university/industry effort, formed under the leadership of the Advanced Systems Institute of British Columbia (ASI), to carry out a Knowledge-Based System (KBS) project in conjunction with MacMillan Bloedel, British Columbia's largest forest industry corporation. The project aims to build a KBS for detecting and predicting faults in large off-highway logging trucks and planning allocation of resources to improve the availability of this equipment. A major catalyst for the development of this type of technology is the Canadian Space Station Program (CSSP), which is anticipated to require extensive expertise in AI technology for development of the Mobile Servicing System (MSS).

MacMillan Bloedel has installed devices on some of its trucks to record 25 types of data at regular intervals while the equipment is in operation. The STAR project will develop a KBS to analyze this collected data, detect performance degradation, predict impending fault states, and diagnose the cause of detected faults. The KBS will make qualitative judgements on the seriousness of actual or anticipated problems. The project is also investigating

the possibility of planning preventive maintenance work to maximize the availability of the equipment.

In order to perform this function, the KBS will be given observations collected by either a human user or the automated monitoring system. It will then monitor these observations for symptoms, form hypotheses about the probable causes of symptoms, predict their near and long term effects, and recommend possible courses of action for alleviating them. The system's reasoning will be based on an internal model of the structure and function of the device and knowledge of relevant diagnostic, repair and planning strategies.

Participants: Timothy Bult, Max Krause, Randy Roesler, Daphne Hassner, Peter George (MacDonald Dettwiler), Richard Rosenberg (UBC), Frits Swinkels (SFU), Eric Manning (U. of Victoria), Brent Sauder (MacMillan Bloedel), Dennis Schwab (Finning Caterpillar).

Knowledge Acquisition Research and Development at Acquired Intelligence Inc.

by Brian A. Schaefer

Recherche et développement en acquisition de connaissances chez Acquired Intelligence Inc.

RÉSUMÉ: Acquired Intelligence Incorporated (AI Inc.) est une corporation canadienne à 100%. AI Inc. s'occupe de développement (outils pour les systèmes de connaissances ainsi qu'applications commerciales) et de services tel que la consultation, les sommaires de recherche, les études de practicabilité, l'éducation, et le développement sous contrat de systèmes prototypes et commerciaux. L'acquisition de connaissances est un thème central à tout cela. Chez AI Inc. l'on reconnait trois sources principales de connaissances, à savoir, les experts humains, les bases de données de cas et d'exemples, et le matériel écrit tel que les livres de référence et les manuels. Leur programme de recherche comprend des projets dans chacun de

Acquired Intelligence Inc (AI Inc.) is a private, 100% Canadian-owned corporation involved in the commercialization of artificial intelligence. Acquired Intelligence's offices and laboratory are located at the University of Victoria, as guests of the Faculty of Engineering, so as to facilitate university-industry collaboration. This has proven to be very successful. Acquired Intelligence is now expanding its facilities to include off-campus offices.

AI Inc. was formed to exploit the widespread commercial need for knowledge-based system products and services. The company is in the process of developing two lines of knowledge-based system products to meet this market demand including:

- Knowledge-based system tools running on affordable computers that can be used directly by experts to create their own expert system applications
- Knowledge-based system applications that provide solutions to known problems and can be marketed to well-targeted business and professional communities.

AI Inc. also provides a wide range of services related to knowledge-based system technology including consulting, literature surveys, educational services, feasibility studies, and contracted development of prototype and commercial systems.

AI Inc.'s staff includes specialists trained at Oxford, Queen's, Victoria and Carnegie-Mellon Universities. The group is highly experienced in project management and the development of artificial intelligence technology. Personnel have experience in the development of knowledge-based system shells, knowledge acquisition systems, database packages and machine learning modules; and they have worked in application domains as diverse as neuropsychology, haematology, commodities and mineral exploration.

An area of research that is central to all aspects of our business is knowledge acquisition. It is our belief that widespread proliferation of this technology awaits the emergence of economical, knowledge-based system tools running on affordable computers, enabling experts to develop and maintain their own knowledge bases.

At Acquired Intelligence Inc. we see three major sources of expertise and knowledge to be acquired for exploitation in knowledge based systems: human experts; databases of cases and/or examples; and written materials such as reference books, manuals, and handbooks. We have organized our research program such that projects are under way in each of these areas. The following will provide a brief description of our research projects in knowledge acquisition.

AI Inc. Knowledge Acquisition System

Acquired Intelligence has developed a proprietary knowledge acquisition methodology. This methodology is designed for the direct elicitation of domain knowledge from human experts. The aim of this project is the development of an automated system for the elicitation of domain knowledge from human experts.

In this project knowledge acquisition has been viewed primarily as a human problem rather than a computing problem. Consequently we have focused on the cognitive psychology literature for insights into solutions to this problem.

This project is specifically concerned with the following issues:

- development of a knowledge representation scheme that is easy for experts to use and is designed from knowledge acquisition considerations
- development of knowledge acquisition methods that facilitate thoroughness and completeness of the knowledge base
- symbolic reasoning with uncertainty (the

Brian Schaefer is the president of Acquired Intelligence. He holds a D.Phil. from Oxford in cognitive psychology.

psychological literature argues against the use of numeric methods for handling uncertainty.)

- development of methods for guiding the expert through the knowledge elicitation process that capitalize on natural pattern recognition skills to facilitate the capture of implicit knowledge.

Acquired Intelligence is currently working on the C implementation of this knowledge acquisition methodology for IBM PC/AT and compatible machines. This system will enable the domain expert to develop simple or complex knowledge bases without the assistance of knowledge engineers.

Knowledge Acquisition System Interface to NEXPERT

Acquired Intelligence aims to make the Knowledge Acquisition System available to users of the most successful commercial shells that run on personal computers. Acquired Intelligence has recently begun investigating the feasibility of developing an interface between our Knowledge Acquisition System and the NEXPERT expert system shell from Neuron Data (this shell is available in Canada through Applied AI in Kanata). This interface will expedite the development of knowledge bases for execution in NEXPERT.

Project SAVANT

Project SAVANT was a very large project aimed at the automation of a knowledge acquisition methodology, the development of an expert system shell for that automated knowledge acquisition methodology and the development of an application in clinical neuropsychology using the tools developed. The resultant system included database facilities and a machine learning module (C-DUCE) based on the DUCE system for structured induction developed by Dr. Stephen Muggleton at the U. of Edinburgh and the Turing Institute. In conjunction with Dr. Russell, a registered clinical neuropsychologist, knowledge bases were built using both the automated knowledge acquisition system and the machine learning system.

This work was carried out primarily on Xerox workstations. The project received technical direction and management from the principals of Acquired Intelligence. Acquired Intelligence has recently purchased all rights to this technology and is currently preparing plans to incorporate aspects of this technology into its ongoing research and development efforts.

Knowledge Acquisition-Machine Learning Environment

Our experience in building automated systems for eliciting knowledge from human experts and our experience with the C-DUCE machine learning approach to knowledge acquisition has led us to embark on a project to combine the two approaches into a single environment. In this environment the expert will be able to move back and forth between the utilities provided by each separate approach so that resultant knowledge bases will be a composite of both expert generated knowledge and machine generated knowledge. This should expand the applicability of expert systems technology to those problems that cannot be handled adequately by either approach alone. Additionally, the environment will allow the construction of knowledge bases using either approach to the exclusion of the other.

An environment will be developed using the Acquired Intelligence knowledge acquisition system, some of the most useful utilities from the C-DUCE system and the SAVANT system for structured induction.

Knowledge Acquisition from Text

The acquisition of knowledge from written materials is routinely performed by people. People consult regulations, handbooks and reference materials in many areas of employment. Acquired Intelligence is in the preliminary stages of a research project, in conjunction with members of U. of Victoria's Faculty of Linguistics, to develop a system that can construct knowledge bases from analysis of written materials. This project will focus on the acquisition of knowledge from books of regulations such as those produced by government departments and agencies. The resultant system could be an enormous aid to those that monitor compliance with regulations as well as those whose sphere of employment is guided by regulations (e.g., the construction and transportation industries).

Applications Testing

Thus far we have been able to test ideas embodied in our knowledge elicitation methodology through the development of a variety of knowledge bases: diagnosis of brain dysfunction in children (Dr. D. Russell, Ottawa); diagnosis of learning disabilities and brain dysfunction in children (Dr. D. Russell, Victoria); haematology — diagnosis of easy bleeding, excessive bleeding, prolonged bleeding and unwanted bleeding (Dr. Jon Gerrard, Winnipeg); interpretation of geophysical instrument readings for mineral exploration (Mr. Rockel and Mr. Saleken, Vancouver); and commodity market analysis (Mr. Bastin, London, U. K.).

Research Support Acknowledgement

For the above mentioned projects Acquired Intelligence would like to acknowledge the support of the National Research Council industrial research assistance program (IRAP), the Science Council of B.C. assistance for applied research program (AGAR), and the Science Council of B.C. graduate research, engineering, and technology scholarship program (GREAT).

For further information on Acquired Intelligence Inc write to Acquired Intelligence Inc., c/o Faculty of Engineering, University of Victoria, P. O. Box 1700, Victoria, B.C., V8W 2Y2, Canada. Phone: (604) 721-8693 or 721-8694, FAX (604) 721-8653.

L'échéance pour le numéro d'octobre est le 15 août

Social Issues Conference

by Robin Cohen

Une conférence sur les questions sociales

RÉSUMÉ: Une conférence au sujet des "Professionnels et de la responsabilité sociale" a eu lieu à l'université de Waterloo du 16 au 18 mars. La conférence, qui était subventionnée par le Centre pour la société, la technologie, et les valeurs, regroupait des gens de plusieurs professions venus discuter de leurs responsabilités vis-à-vis les codes d'éthique, la paix globale, et les groupes d'activistes. L'auteur de cet article s'est attardée au travail en informatique.

16–18 March 1988, Waterloo, Ontario

This conference, sponsored by the Centre for Society, Technology and Values at the University of Waterloo, discussed the topic of professionals and social responsibility. The conference brought together about 100 individuals. Intended to get people from various professional organizations discussing the social and ethical responsibility of each profession, the title ("Professionals and Social Responsibility: Conflict or Congruence?") was designed to examine whether various professions have common goals in assuming social responsibilities, or whether there are sufficient differences among the professions, to prevent useful dialogue toward common goals.

The keynote address on the evening of the 16th was by Jack Stevenson, of the Philosophy department at the University of Toronto. Stevenson's main concern was to provide some theoretical grounds for "reasonableness", to provide professionals with some grounding for achieving social responsibility. He commented on how pure scientists often divorce themselves from these concerns, and how applied scientists face a complexity of interests — between themselves as professionals and their clients, and between the interests of their clients and the interests of society. He then commented on scientists' usual concern for rationality as the main yardstick of successful progress. But certain actions of professionals, while beyond rational criticism, still require scrutiny. Stevenson then suggested replacing the desire for rationality with a desire for "reasonableness". This then requires a formal concept of reason, together with interpersonal dialogue and experiences of living. This moral dialogue, injected with "caring", can take scientists beyond reason toward social responsibility.

It was interesting that this talk which could be interpreted at one level as a detailed philosophy lecture, generated questions only from the computer science people in the audience (three in AI, one in graphics). Some of the concerns and proposed solutions evidently struck a chord with these computer scientists.

The first morning of the conference was devoted to the topic of codes of ethics, in various professional groups. Mark Frankel of the American Association for the Advancement of Science's office of scientific freedom and social responsibility delivered the message that professionals must feel the responsibility to report misdeeds in their profession. This is best achieved with some professional self-regulation, as in having a code which is highly visible. Obligation to codes could then be brought to legislation (rather than the counterproductive procedure of having mandatory laws to actually control the actions of professionals).

Steven Unger from Columbia University specializes in ethics for engineers. He warned that codes of ethics cannot embody a "world view" and that ultimately different individuals must reach personal decisions about what to do, ethically. But codes should universally support the desire of professionals to get their job done right and to get the right job done. This has been overlooked in some cases. For instance, Unger discussed some specific cases of engineers who reported misdeeds and ended up being sent to distant locations and eventually fired.

Leonard Brooks, from the Faculty of Management at the University of Toronto, focused on ethical standards within business. He mentioned the growing concern of corporations to ensure ethical practices — to maintain the confidence of clients. Interestingly, in Canada the concerns of corporations for ethics among employees is high, but towards customers and suppliers is much less than world-wide interests. Currently, most of the ethical standards are set separately by each company.

Subsequent to these talks on codes of ethics for professionals and how to implement them, a panel, chaired by Conrad Brunk of the University of Waterloo, investigated further some of the issues raised. Abbyann Lynch, of London's Westminster Institute for Ethics and Human Values, raised the question of what we consider a "professional", and why people are concerned with ethics. If it is just a tool, we could end up trivializing ethics.

Arthur Schaefer of the University of Manitoba is regularly called in by professionals to mediate ethical discussions. He mentioned several cases of mediation for the medical profession, where he personally argued that patients be involved in the decision making process and was consistently told that this was not in their best interests. He then raised the question, introduced earlier by Stevenson, of whether his responsibility was more to his clients or to the community at large.

Schaefer also had interesting views on whistle-blowing, as mentioned by Frankel. Schaefer suggested

Robin Cohen is a professor in the Department of Computer Science at the University of Waterloo.

that professionals must sort out their transgressions, rather than going public — to admit that even professionals are fallible and to keep the image of the professional as competent.

After the first half day, the issues at hand seemed to be how professionals sort out:

- personal values
- commitment to clients vs. third parties
- how to go beyond science to reasonableness
- whether codes of ethics should be implemented with:
  — greater push for reporting and correcting, possibly going public, or
  — greater self-regulation, and concessions to possible mistakes.

The afternoon of the first day shifted topics to the major concern of several professional organizations — global peace — and the related topic of human rights. Human rights was addressed by Rose Sheinin, a biochemist at the University of Toronto, who commented on problems faced by women professionals, to advance and be accepted in their chosen professions. Examples from the academic environment were presented, to show how few women have advanced to positions of authority within academia. Sheinin also argued that girls have a general lack of role models throughout education to encourage progression into the scientific disciplines.

Human rights was also addressed by Israel Halperin of the Canadian Committee of Scientists and Scholars, who discussed his personal efforts to free particular scientists suffering oppression in various countries. He commented that his campaigns are merely one contribution to the eventual success of freeing these scientists, and are intentionally devoid of particular political interests.

Anatol Rappaport, a professor of Peace Studies at the University of Toronto, then shifted to the general topic of the source of redemption for scientists. He argued that science cannot be entirely value-free, continuing points raised by Stevenson about the moral nihilism most pure scientists seek. Rappaport's main concern is about the abuse of science in aid of the military. Scientists then have an individual responsibility to refuse to contribute towards the war machine.

This includes the responsibility of scientists who are employed by the military, and the responsibility of scientists to study peace, scientifically. This latter aim, though constructive, is hard to achieve, since "social science" is harder to defend than pure science. Computer science people in particular can play a role, according to Rappaport, in exposing the infeasibility of the Star Wars project. But beyond this, we should make a positive effort to speak out against war, regardless of feasibility arguments.

The challenge of global peace was investigated further by a panel chaired by Tom Perry of Canadian Physicians for the Prevention of Nuclear War, including psychiatrist Joanna Santa Barbara, Dr. Gilles Hurteau and social worker (and Hiroshima survivor) Setsuko Thurlow. Perry made clear the connection of global peace to concerns of the professions. For doctors at least, the whole purpose of the profession — to prolong life — is made pointless in the event of a nuclear war. Moreover, it is important for doctors to speak out against the ratio of funding allocated to health care, relative to the money invested for war.

Panelists suggested educating children to adapt to the spectre of nuclear war, by fostering positive, problem-solving attitudes, and responsibility for the preservation of the planet. The educational system needs to instill critical thinking, giving children the truth, together with a sense of responsibility. Some efforts have been made by boards of education but these have been far too sporadic to date.

In the discussion that followed, people commented on the need to humanize science — not to give up on its rigour, when teaching, but to educate scientists to have a human element. There was an interesting concern that some professionals may be overstepping their realm — taking on problems which properly belong to some other professional's domain.

In summary, to assure that professionals act responsibly, we must be particularly concerned about:

- equal access within the professions
- freedom for professional duties, in all nations
- special concerns which beg the existence of particular professions, like global peace.

Mental health and social work professionals can play a role in preparing future generations for these difficult moral decisions, and all educators bear some responsibility for making students aware of future responsibilities.

Friday morning the conference shifted to the topic of activist groups. Beth Savan of Environmental Studies at the University of Toronto delivered a key address on the need for professionals to take on social concerns.

She mentioned some of the problems with professions as they are organized today. A hierarchical structure — say of senior scientists with labs, served by junior scientists who do the leg work, down to the everyday work of students — relies too heavily on the goodwill of a mentor for advancement. Moreover, this is a stressful training ground (e.g. for doctors), which ends up being counterproductive for all concerned, especially the public (e.g. overworked interns getting callous to the needs of patients).

Some activist groups try to break this mold, according to Savan, including support for individual professionals who want to pursue socially beneficial research. The appearance of these groups can undermine the general view of the profession to society, bringing pressure to bear on the larger group to also act responsibly. Savan, in fact, advocates that all professions include commitment to the wider society in their practice.

Questions after the talk included concern that professionals would abuse their position for personal aims. This continued well-presented points by Unger that individuals must follow their own values. Savan admitted some abuse was inevitable, but the larger benefits still made the approach worth pushing.

There was also concern about involving the public in decisions regarding the direction for professions. Savan pointed out that directions of professionals are already largely dictated by corporations, through funding, who are not any more educated in making the right decisions for the profession.

Continuing earlier discussion of advantage for professional groups to self-regulate, Savan mentioned that groups whose clientele is a powerful elite group, do serve the community better by not necessarily relying on the client to dictate the directions of the group.

Following Savan, activist group representatives discussed particular efforts, including Margaret Keating and Diana Dick of the nursing profession, Gary Chapman of the Computer Professionals for Social

Responsibility (CPSR), Karen Messing of the Groupe de recherche-action en biologie du travail, at the University of Quebec at Montreal, Michael Rachlis of the Medical Reform Group of Ontario and Steven Shrybman of the Canadian Environmental Law Association.

The nursing representatives discussed issues of particular concern to their group, including the allocation of health care resources, and urged greater involvement in policy development, through involvement in politics and use of the popular media.

Messing, involved with occupational health and safety, cited several cases where the efforts of her group assisted workers who would otherwise not have been treated fairly — including possible discrimination against women. She gave some strong reasons for involving the public (those affected) in the decisions of the professionals in these cases. There is still a problem with the public being sceptical of the "experts", but the effort to keep the relation equal can be pursued.

Rachlis commented on efforts of the medical reform group to take stands on particular issues — the drug bill, women's choice in abortions, mid-wifery. Shrybman commented on some of the ethical responsibilities of lawyers, citing efforts of particular lawyers to reach beyond corporate clients to those truly in need.

A few sparks got set that morning as well when Gary Chapman of the CPSR decided to address some issues raised earlier at the conference, rather than focusing on his prepared talk about computer professionals and related issues.

He presented some controversial arguments against the proposals for global peace, educating children to try to save the planet. He felt that "caring and sharing" is not the reality of today. Moreover, he felt that freely associating with Soviets was not productive. He stated that CPSR's position is that there be no official contact between American and Soviet computer scientists, until an independent Soviet route is established towards global peace. It seemed that his radical position was that the only way to assure that the Soviets will do their share of the contribution (towards peace, which is a critically valuable goal) is to refuse to have a dialogue on these matters until they work just as hard, independently, for the same aims. And Chapman felt (brought out in conversation which continued over the lunch hour, with conference participants grouped in a circle at the buffet counter) that Soviet scientists often just rode on the coattails of the efforts of the American scientists.

Chapman thus presented an alarmist view — warning that while people sat back pushing for global peace some real opportunities to use military force to turn around critical social problems — e.g. South Africa — would be overlooked.

In discussion at lunch-time, several people still pressed for dialogue through to proper solutions. The basic differences between Soviet society and ours were offered as reasons for expecting less of the Soviet scientists at this point.

Friday afternoon concluded with one more talk about a particular group, the association of professional engineers of Ontario, represented by Jeffrey Tyndall. He discussed difficulty in interesting his community in instituting proper procedures for whistle-blowing, reflecting concern in general about the degree to which engineers are concerned about social responsibility.

Then, Kelly Gotlieb, professor of computer science at the University of Toronto and Robin Cohen professor of computer science at the University of Waterloo commented on some efforts to educate future computer science professionals, through courses offered at the two universities. The University of Toronto course, open to all students with at least one computer science course, has been offered for a number of years. Students are encouraged to think about social issues from computerization, through course credits such as debates, a media study and a personal record project — trying to get hold of a record held on them and reporting on the incident.

Robin Cohen commented on the challenges of directing a course for fourth year computer science undergraduates, taking advantage of their computer science knowledge to assume understanding of various technologies, to then study particular social effects. Students are eased into course projects by lectures which discuss how to make ethical judgments and do social science research, a visit by library representatives, an early exercise in writing (a book report) and feedback on a project presentation, delivered to the class prior to the completion of the project.

At Waterloo, computer science students can supplement their education in this area with a general course in society, technology and values, STV100, open to students in all faculties and team taught by professors with varied expertise. In addition, the Centre for Society, Technology and Values sponsors a colloquium series on the subject.

Leonard Waks of Science, Technology and Society program at Penn State University, commented on efforts in general, in the United States and elsewhere, to educate students about science, technology and society. He stressed the need for a concerted effort to bring about this education, including involvement at the high school level.

In summary:

- professionals should look toward breaking the limits of the hierarchical structures within their professions to embrace greater concern for the public
- efforts to introduce this awareness are not without their problems
- activists still feel they are increasing public awareness
- education can continue to play a role.

To close, Willem Vandenburg of the University of Toronto's Centre for Technology and Social Development, did a conference wrap-up, commenting on shared patterns of social responsibility in the various professional groups. He also felt that a greater "macro level" realization was required, to assure success of all the "micro-level" projects. More conferences could keep us communicating.

In general, the conference seemed to present more concerns than solutions. Some challenges were raised for professionals interested in acting responsibly to the public, both as individuals and for their professions as a whole. For computer scientists in particular, academics can play an important role in educating, researchers can consider the possible applications of their research, business people can tread the ground between responsibility to the employer, the client and the general public. But scientists need not be, perhaps should not be, value free.

Proceedings for this conference will be compiled shortly; the conference organizers are currently investigating the best avenue for publication. Inquiries may be directed to: the Centre for Society, Technology and Values, University of Waterloo, Waterloo, Ontario N2L 3G1.

Report on the 1988 Distributed Artificial Intelligence Workshop

by Ernest Chang

Rapport au sujet de l'atelier 1988 en intelligence artificielle répartie

RÉSUMÉ: L'atelier 1988 en Intelligence Artificielle Répartie (DAI) a eu lieu au centre des conférences de l'UCLA à Lake Arrowhead en Californie. Il y avait 35 participants et plus de 120 articles soumis. Les sujets abordés comprenaient la théorie des agents, la théorie de coordination, les implantations et la DAI en général. Les discussions ont également essayé de clarifier le concept de DAI.

22 - 25 May 1988, Lake Arrowhead, California

The UCLA Conference Center is located in the heart of the San Bernardino Mountains at Lake Arrowhead, five thousand feet above the smoggy plains of Los Angeles. The grounds of the Conference Centre are well-groomed with a putting green, tennis courts, a swimming pool with indoor jacuzzis, and gardens of irises to set off the main lodge done up in gray cedar shingles and gables. The guests stay in cottages nestled in the ponderosa pine groves among manzanita blooms and lupines. Bells announce meals served at the immaculately set tables, with the highest standards of cuisine. All this was the setting for the 1988 Distributed Artificial Intelligence (DAI) Workshop.

The tradition of this workshop has been attendance by invitation only, but this year a call for papers was issued. The process yielded a total of 35 participants and observers, who were privileged to attend the three day session lasting from Sunday night, May 22, to Wednesday noon. Presentations took place during the mornings and evenings, so that most afternoons were free for informal discussion and environmental exploration.

Distributed Artificial Intelligence has been a name associated with a loosely connected body of work, looking for a definition. Carl Hewitt described the distinction between artificial intelligence and traditional computer science in the following way: computer science studies algorithms that avoid inconsistency while AI deals with inconsistency directly. The difference between distributed computing and DAI is that the former studies algorithms operating in a network, whereas DAI assumes the existence of autonomous agents (that can deal with inconsistency), and studies the principles and techniques for coordinating interacting agents.

Much of the general discussion addressed the broad problem of characterizing the domain of DAI. Sridharan, now at FMC, reviewed the proposals from the last DAI Workshop. For the DAI system as a whole, consider: number of agents, the organizational structure being modeled (committee, panel, institution), the collective sense of purpose, control strategy, global properties explicit or emergent, homogeneity of agent type, system dynamics (fixed to self-organizing behaviour).

For the agents, consider the following dimensions: agent granularity (resemblance of single agent task to the collective task), degree of autonomy, agent dynamics (fixed vs adaptive behaviour), agent coupling (how well does one agent model another), agent resources, extent of agent interaction, agent intelligence. In essence, according to Sridharan, the problem is really that of DAI or distributed intelligent agents, rather than of distributed AI.

Toward the end of the workshop, Miro Banda of Boeing AI Labs put forward the tetrahedral model of DAI, that formed by the interaction of the poles DAI, coordination protocols, execution platforms (machines), and agent theories. From this model came the formulation of six fundamental problems: creation of the organizational structure maximizing the intelligence of the DAI system, description of coordinating knowledge that agents need to know, studies of distributed processing & processors, negotiation and interaction of agents, tools for building agents, group display technology, and validation and verification of DAI systems.

The presentations fell into four major groupings: those dealing with the theory of agents, those dealing with the theory of coordination, those dealing with implementation issues, and those of general interest. This article is not a precis of each talk in the workshop but selections from a few in each area that were personal highlights for the author. Apologies are made for all omissions, errors and misinterpretations of the views put forward by the workshop participants.

Theory of Agents

In the theory of agents are included the talks by Grosz & Sidner, Cohen and Hewitt & Star. The approach taken by Barbara Grosz of Harvard & Candice Sidner of BBN is to understand some of the goals of an agent through analysis of two-person discourse processing. Specifically, the performance of a cooperative task requires collaboration in its planning and action. The overall plan is more than a union of the plans for the actions of each individual; rather, each person has individual plans that are not blueprints for action but are instead incremental

Ernest Chang is the Technical Director of the Advanced Technologies Department, Alberta Research Council, Calgary. His primary responsibilities are to guide the department's research activities in the areas of AI and automation.

recipes, collections of actions achieving some goal, which are dynamically reassessed during action. Furthermore, both agents need mutual beliefs about the other's ability to execute, need intentions about actions they would do to contribute to the overall plan, and need beliefs about the other's intentions. In collaborative planning, agents must negotiate goals, negotiate the recipes that they might use, and negotiate the agency chosen to execute the actions.

The Speech Act contribution to this DAI Workshop was given as usual by Phil Cohen of SRI, who discussed the theory of joint intention based on a theory of commitment. Briefly, this talk attempts to relax Searle's (the father of Speech Act) notion that collective intentions are primitives. Thus, assume that a *commitment* to a goal G means that the choice of G is persistent over time, but not necessarily forever, and that an agent will attempt to achieve G by making it true later, it being false now. An *intention* to do a goal G is to commit to doing it successfully and to believe that it will happen, the commitment being dropped only if the goal is satisfied, or the agent comes to believe that it will never be satisfied, or the reasons for having the intention have changed. A *joint persistent* goal P is one that both parties believe to be false initially, that they both commit to making true, and that they both commit to until they either both know it true, both believe it will never be true, or both agree that the reasons have changed. In a *joint intention* to do A, both parties are committed to *our* having done A, mutually believing we were about to do A. Thus, in joint actions in which X and Y perform two acts in sequence, one first derives the requirement of synchronizing actions, and then obtains that joint commitment to the sequence which means that if X acts first but Y fails, then X's commitment to the joint sequence will cause X to start over. In conclusion, joint intentions require synchronization for starting a sequence, communications to agree on joint intentions, and signals for termination since one party alone cannot terminate a joint agreement.

Leigh Star of U. of California at Irvine & Carl Hewitt from MIT proposed new units of analysis for distributed AI that will apply to humans as well as machines in open systems. The most important unit is that of commitment, which deals with the use or provisioning of resources. When two or more commitments cannot be met, then a trial of strength rearranges the commitments, and synthesizes new ones. Thus, relationships between commitments and actions are based on trials of strength. From the notion of commitments, they assert that beliefs and intentions arise when you have commitments about commitments, that cooperation is a situation in which parties have become committed to each others' commitments, and that robustness means the ability to keep commitments in the face of trials of strength.

Theory of Coordination

The talks on theories of coordination mostly dealt with the subject of negotiation. Ed Durfee, from U. of Massachusetts at Amherst spoke of negotiations for the allocation of network resources. The framework for negotiations requires: a protocol, which is a structure for communications called a Partial Global Plan; common knowledge (agent A knows that agent B knows that agent A knows...) about the meta-level organization of the system; and finally a reasoning component. Partial Global Plans consist of information about objectives, plan-activity-maps, solution construction graphs, and status information.

For Katia Sycara of Carnegie-Mellon, negotiation is an iterative process involving utility or payoff functions. Persuasive arguments change the payoffs of a proposition either by changing its importance to an issue in terms of unpleasant consequences, or by showing that alternatives exist, or by demonstrating that the proposition furthers another important goal. Generating a persuasive argument requires a belief-tree, which is a causal network leading to a goal.

In negotiation among cooperating experts, Susan Lander of U. of Massachusetts at Amherst assumes that the experts have conflicting local goals, and that there may be globally inconsistent knowledge but no central authority. Given overlapping but inconsistent solutions based on different areas of expertise and problem-solving strategies, conflict resolution is necessary. This can be done either by compromise, in which one or more agents relax or relinquish goals, by reformulating goals, or by third party intervention. A model of compromise bargaining involves an iterative cycles of solution proposals, evaluations and revisions.

Implementations

The talks that I group into "implementations" include presentations by Cindy West, Mark Fox, Lee Erman, David Lane, Les Gasser and Richard Korf. The project that Cindy West of U. of California at Davis described is the implementation of an assumption based truth maintenance system using multiple agents, in an environment where information from other agents cannot always be trusted. The application involves multiple seismic stations, each modeled as an intelligent agent. The result received from another agent would only cause a station to process local information to look for additional features, but never to revise a local belief except where local results warrant the change.

Mark Fox from the Robotics Institute at Carnegie-Mellon described a method of using knowledge of constraints to guide the solution to a distributed problem. The central thesis is that such knowledge can be used to develop an approximate topology of the problem space, with the job shop scheduling problem as an application. Specifically, micro-opportunistic constraint-directed search (CDS) is the most promising approach, using aggregate planning at higher levels and problem space textures at the local level. Aggregate planning is based on a statistical view of expected demands from job agents, either in the immediate time horizon or as forecast, and is done at the tool level where committed and anticipated demands are maintained. At the local level, problem space textures are measures of the nature and difficulty of decision islands, those subgraphs of directly interacting objects (resources) and constraints. Textures are described by looseness, contention, elasticity, importance, connectivity and reliance of constraints and resources. For example, temporal elasticity can be described, using James Allen's temporal relations, as the range of intervals within which each activity can occur. At the local level, each job agent would generate bids or sells for resources, and a moderator would use negotiation operators to produce an optimal set of resources allocated to all jobs as a whole.

Module Oriented Programming (MOP) was described by Lee Erman of Teknowledge. The basic computational unit of a system is a module, either primitive or composite, which either produces or consumes resources, and always has an embedded manager. The notion is

that process is separate from that of module, so that processes can be assigned to modules for execution. This model allows a natural implementation of agents for DAI systems.

The workshop coordinator, Les Gasser of USC, gave a model of distributed problem-solving requiring a definition of the social structure and coordination framework. These structures may differ from different perspectives, and a representation scheme would be required to reflect these different viewpoints and any reconfigurations needed to adapt to new situations. He then went on to describe a frame-based implementation in which problems are modeled as a mixture of settled and unsettled sub-problems, the settled ones being the expected behaviour of others. In this system, the assignments of roles for any coordination framework are supported by the the explicit description of agents and their actions.

David Lane and Mike Chantler from Heriot-Watt University, Edinburgh, described a method for dealing with the control of autonomous undersea vehicles as a DAI system. In particular, the vision sub-system can be seen as a hierarchy from low-level filtering and segmentation to feature extraction, anomaly detection to classification, with a blackboard system as the natural implementation. The distribution approach used is a vertical one, in which each agent would process and reason about a part of the image, exchanging partial image results across a level as needed. This is claimed to be better for load balancing.

An approach to DAI using generalized game trees was given by Richard Korf of UCLA. A chess game can be seen as a multiple autonomous tightly coupled system between two agents, with the machine agent capable of acting at an expert level. The n-agent game tree is extendable from the two-agent one, and the general notion is that an opponent would model the strategies of other players, and their models of the strategies of others' strategies, etc. Communication would permit more exact evaluation of others' strategies, and flattens the knowledge hierarchy in the sense that the model of another's strategy can simply be the other's strategy.

General Interest

Among other talks of general interest is that of Thomas Malone from MIT. He described a new joint initiative among computer scientists and social and behavioural scientists called Coordination Theory. This is motivated by the movement to computer-supported work, by the organizational changes in human agencies. These were in turn made possible, by the improvements in parallel and distributed computing systems, and by the possibility of theoretical convergence among many disciplines. Tentatively, coordination is described as the additional information processing needed when multiple connected actors pursue goals that a single actor pursuing the same goals would not perform.

Another interesting talk was given by Murray Glazer of the University of Toronto who presented a DAI in terms of knowledge. After giving a formalism for computation and process, group and common knowledge, he showed that commitment, negotiation and blackboard systems can all be understood in terms of what each agent knows and when the information is known. The point is that using knowledge requirements, it is possible to understand what properties a system needs to satisfy the desired characteristics in terms of coordination and coherence.

Conclusion

The general feeling of the workshop was that DAI is an area which involves many disciplines, that it is a growing area of interest, and that for the workshops more time should be given to study sessions and brain-storming in addition to formal presentations. The next DAI Workshop is scheduled to be held in the Seattle area, with Miro Banda of Boeing AI as the organizer.

The Fourth IEEE Conference on Artificial Intelligence Applications

by Betty Ann Snyder

La quatrième conférence de l'IEEE sur les applications de l'intelligence artificielle

RÉSUMÉ: Cette conférence, tenue à San Diego, a attiré environ 800 personnes. Elle fut précédée de deux jours de cours intensifs sur divers domaines d'application de l'AI. Les sujets traités à la conférence inclurent le diagnostic, la vision et la robotique, le dessein, la fabrication et le contrôle de processus, et les applications de planification. Les conférenciers invités furent Scott Flaig (DEC), Scott Fahlman (CMU), et Henry Lum (NASA).

14 - 18 March 1988, San Diego, California

Imagine, if you will, the task of summarizing a conference that spans almost every topic known to man. Then imagine assigning the job to a natural procrastinator.

Well, as you have by now deduced, it is the day before the confounded article is due. And frankly, at the moment I would rather write a summary of the acrobatic interaction of the swallows currently entertaining me, than a summary of the Fourth IEEE Conference on Artificial Intelligence Applications.

My only major complaint about this conference concerns the location. The San Diego location of the Sheraton Harbor Island Hotel is not conducive to concentration. However, despite the proximity to beaches, the San Diego Zoo, Seaworld, and La Jolla, about 800 people battled their hedonistic tendencies to

Betty Ann Snyder is a contract specialist at the Alberta Research Council in Calgary, Alberta.

attend selected sessions. The conference was sponsored by the Computer Society of the IEEE, in cooperation with the AAAI, so it attracted a good selection of submissions. James Miller, the "General Chair", indicated that the paper submissions were up 50 percent over last year's conference.

A list of all conference submissions would be extremely useful for an analysis of the hottest areas of AI research and development. Unfortunately, we do not have a way of gathering all this information. The "Program Chairs", Elaine Kant and Dennis O'Neill, note that there was "substantial participation by overseas colleagues" and that there were more submissions from industry than from academia. The degree of participation from industry is, of course, always hard to gauge due to the existence of proprietary artificial intelligence applications.

The conference was preceded by two days of tutorial sessions. The subject choices were as follows:

- Knowledge System Development *A. Barr*
- AI in Manufacturing *M. Fox*
- Knowledge Structuring for Knowledge Acquisition *E. Miekle, M. Garrison*
- User Interfaces *M. Stelzner, A. Cypher*
- Natural Language *G. Hendrix*
- Object-Oriented Concepts *K. Schmucker*
- Doing it on a Mainframe *J. Aikins, P. Harmon*
- Survey of DOD Research and Applications in AI *Bob Simpson*

One dark note in the tutorial schedule was the absence of Paul Harmon. Jan Aikins presented this tutorial at short notice.

Keynote Addresses

Impact of AI on the Corporate Enterprise, with Emphasis on Computer Integrated Manufacturing

The talk given by Scott Flaig of DEC on "Using Technology as a Competitive Edge" summarized some of the 12 years of AI development at DEC. Digital is now using AI techniques for large training applications, such as job matching, apprenticeship, technology transfer, and the training of over 300 people in manufacturing skills. The XCON project for system configuration has been further developed to include about 10,000 rules. Expert systems are being developed for diagnostic purposes, process evaluation and process control. Manufacturing applications of the technology include scheduling, materials handling and inventory control. Expert systems are also being implemented for distribution, carrier managing, a dispatching network system and distributed diagnostics at the customer site. XSEL assists with sales and equipment sourcing. Their newest applications involve prediction of markets, modeling and simulation and intelligent decision support. DEC now has 150 staff experienced in artificial intelligence/expert system applications.

Scott indicated that integration of cross functions, and integration with conventional technology hardware and software will play a larger role as expert systems are brought together with various existing applications.

What's Missing in AI? — Can Massively Parallel Architectures Help?

Scott Fahlman of CMU addressed the issue of building systems capable of handling broad based knowledge fields, complex multi-level recognition problems and non-symbolic representation and learning. Currently, if knowledge is not well defined within a narrow area, we are unable to efficiently deal with it to build an expert system. One solution may be to implement active processing power, or active knowledge which will "answer" to matching questions. In an effort to provide active parallelism to solve this problem, CMU is currently researching the use of a parallel network of simple processing elements, wired to neighbors, which is presented with a training set. This network "learns" slowly, but the learning speed is being reduced dramatically by dynamic control of learning step size, and the elimination of "stuck" units which have bogged down. Using these techniques on an IBM RT, CMU has been able to produce learning in 40 presentations, a distinct improvement over the initial requirement of thousands of presentations of the learning set. Further development of this idea will probably require new architectures, better algorithms, and a combination of several existing models to provide variable plasticity and generalization.

Papers

As stated earlier, it is difficult to provide a summary of the incredibly diverse subject areas covered by the conference. The subject areas can be divided into diagnosis, vision and robotics, design, manufacturing and process control, and planning applications. Other topics concerning the mechanisms of system development included program development aids, knowledge representation, diagnostic reasoning, learning systems, tools, implementation algorithms and natural language interfaces. At this point, I will mention the titles or topics of most presentations.

The area of diagnosis was represented by papers on AI-Test for electronic troubleshooting, PERF-EXS for power plant diagnosis, and CONSOLIDATE for machine fault diagnosis.

Vision and robotics were represented by information on a connectionist approach to primitive shape recognition, cooperative focus and stereo ranging, geometric motion planning under unexpected obstacles, truth maintenance in model-based object recognition, and ANDES for knowledge based scene analysis.

Interesting design systems include transformer/inductor design, custom VLSI physical design tools, MES for reuse of transmission equipment, an expert system for bridge design, "Argo" for design by analogy, learning of preference rules for VLSI design, and a knowledge based system for database design.

The manufacturing and process control systems include INCA for process planning in a PCB assembly line, Arthur Anderson's Dynamic Rescheduler, and the CABPRO expert system for process planning of assembled multiwire cables.

The planning applications include a knowledge-based planning/replanning system for Naval command and scheduling, VIPS visual interactive planning system, and an expert system proposed by the University of Genoa for alleviating overloads in electric power systems.

Thursday's invited panel debate on the future of Lisp machines, while not resolving the issue, added colour and intensity to the conference.

In summary, the conference successfully showed the range of artificial intelligence applications currently being developed.

Book Reviews

edited by Graeme Hirst

Visual Reconstruction

Blake, Andrew; and Zisserman, Andrew
[University of Edinburgh]

(The MIT Press series in artificial intelligence)
Cambridge, MA: The MIT Press, 1987, xi+225 pp
Hardbound, ISBN 0-262-02271-0, US$25.00

Reviewed by
Gregory Dudek
University of Toronto

Understanding how the visual process operates, either in the context of human or machine, remains a formidable problem. In particular, drawing a conclusion about what the three-dimensional world looks like from only a two-dimensional projection (on the image) implies "making up" for an entire lost spatial dimension. In fact, there are many hypothetical three-dimensional scenes that could give rise to any single two-dimensional image. One popular hypothesis regarding how the visual system selects a three-dimensional interpretation is that it constructs the one that is most consistent with certain assumptions about the smooth cohesive properties of matter. This is the core of the approach eloquently espoused by Blake and Zisserman.

The recurring theme in many recent approaches to low-level bottom-up vision has been the use of "smoothness constraints"; that is, the preference for world models whose physical properties tend to vary gradually across space and time. Typical examples of such constraints are those on surface depth or derivatives — reasonable surfaces have surface depth values that vary continuously across the object. Such constraints can be elegantly expressed in the language of variational calculus. Differing formalisms for the concept of gradual change thus produce differing classes of surface model.

Blake and Zisserman painlessly introduce the conceptual basis for such constraints and derive the consequent conclusions about the performance of surface fitting algorithms. By relegating the actual solutions of the variational forms to the appendices, they produce a running dialogue outlining the significant conclusions without incurring excessive mathematical overhead. Although the details of the technique require a certain mathematical sophistication, the essential concepts should be accessible to readers with the barest familiarity with the concepts of calculus.

Their approach provides a very readable introduction to the basic concepts (and some shortcomings) of surface reconstruction using "energy" minimization. The penalty functions they use for smoothness are on the first and second spatial derivatives and correspond to surface models that behave like membranes and plates, respectively.

In the later part of the book, the problem of discontinuity detection is dealt with and several approaches to the inherent convexity of this problem are discussed. By the time this discussion culminates in the authors' own "GNC" algorithm — a continuation method for the problem of surface fitting with discontinuities — the level of description has become substantially more mathematically sophisticated than in the early part of the book. The utility of relegating parts of the analysis to the appendices at this stage is rather questionable. A reader comfortable with the later chapters will probably find himself flipping back to the appendices.

Blake and Zisserman manage to introduce the problems of surface reconstruction through functional minimization in a way that is readily digestible and even fascinating for novices in the field. At the same time they present enough detail to make the book interesting to more sophisticated readers. In reconciling the needs of readers across the spectrum of sophistication, they make the book somewhat less than ideal for either. Many readers will be left feeling the discussions are both simplistic in parts as well as too difficult in others. More important, however, is the fact that they succeed in making a potentially difficult subject area accessible to a broad range of readers.

Greg Dudek is a Ph.D. candidate in the Department of Computer Science, University of Toronto. His research concerns visual recognition of objects.

Logics for Artificial Intelligence

Turner, Raymond
[University of Essex]

(Ellis Horwood series in artificial intelligence)
Chichester: Ellis Horwood, 1984, 120 pp
Distributed in Canada by John Wiley, Cdn$43.95

Reviewed by
Lenhart K. Schubert and Francis Jeffry Pelletier
University of Alberta

This short (120 pp) book is intended, apparently, as a supplementary text or reference for a graduate AI course that emphasizes logical approaches to the subject. It gives concise introductions to dynamic logic, many-valued logic, temporal logic, non-monotonic logic, type theory, and fuzzy logic. As can be seen from this diverse list of topics, each is given only a cursory treatment; and to get much out of the author's short discussions, some general sophistication in logic is expected of the reader. (The author's remark, p.16, that his two-page summary of classical logic is all that is required, is seriously misleading.)

Besides the worth of the concise introductions to different areas of logic which an AI student should become aware, the book contains much to recommend it. For one thing, Turner emphasizes the formal semantics of these systems of logic — something which is much too rare in AI texts. In various of the areas he gives good critical assessments of proposals that have been put forth in the literature, and his proposal in Chapter 5 for a semantics of non-monotonic logic based on partial models is especially interesting. And to top it off, the writing is generally quite lucid.

Unfortunately, there are also a number of shortcomings in the book. The main one is its brevity. It is just not possible to introduce classical logic in two pages or modal logic in two pages. It is only the student who already knows the basics of these areas, and of formal semantics generally, who will benefit much from this book. Furthermore, logic is treated in an axiomatic manner — probably the least common form used in AI — rather than in a resolution-based manner or in a natural deduction format. No treatment is accorded to equality, which is certainly central to AI applications. From this it would seem that the main use of the book would be as a guide to the literature of the areas covered. The presentations of the various logics generally stop at definitions with little indication of what the formal properties of the logics are or what they have done or might do for AI. An instructor planning to use the book as a text would have to elaborate considerably on the material. Furthermore, there are large gaps in the coverage. One would think that a book with this title would include a discussion of applications of *standard* logic to AI, such as resolution and answer extraction. And there is no mention of probabilistic logics, planning logics, inductive logics, intensional logics, relevance logics, paraconsistent logics, or applications of any logic to natural language processing. As a text it also falls short by having no exercises. Finally here, we remark on the truly astonishing number of typographical errors in the book, sometimes making text and formulas incomprehensible.

We close with one general remark on Turner's overall division of logics. Like many other authors, Turner divides "non-standard logics" into "rival logics" and "extensions of classical logic". The intent is to distinguish those logics with the same vocabulary as classical logic but with different theorems (usually fewer e.g., many-valued logics, intuitionist logic, relevance logic, fuzzy logic) from those that add new vocabulary and new rules so as to retain all classical theorems and add new ones (e.g., modal logics, intensional logics, temporal logics). Turner notes that the distinction is "not watertight"; but the real situation is much more muddied than that, and Turner should have indicated so. Whether a logic is a subsystem of, an extension of, or a rival to classical logic depends on how we formulate the systems and compare them. Consider, for example, a modal logic which is formulated with "—>" for "strict entailment" (rather than with a box for necessity). Such a system is a subsystem of classical logic in which the —> is interpreted truth-functionally (with appropriate axioms governing it). But now consider formulating classical logic with just {&,~} and the rule of inference: from A and ~(A&~B) conclude B. Given this formulation of classical logic, our modal logic emerges as an extension when we add —> (with appropriate axioms). Similar remarks can be made about many-valued logics (with "assertion operators", as in Turner's Chapter 3), about intuitionistic logics (as in his Chapter 4), or temporal logics (as in his Chapter 2). The alleged distinction between rival ("deviant") logics and extensions becomes vacuous, and only confusion is engendered in suggesting it without a restriction to a specific vocabulary and formulation of the rules.

In sum, we find Turner's book to be most useful as a guide to the literature in certain fields of logic-as-applied-to-AI. It would be more useful in this role if it had broader coverage. It is not suitable as an introduction to the fields themselves because of its lack of depth. And the instructor who attempts to employ it for any use would be wise to prepare errata sheets.

Len Schubert is a Professor in the Department of Computer Science, University of Alberta. Jeff Pelletier is in the Departments of Computer Science and Philosophy at the same university. Their research includes natural language understanding systems, with an emphasis on logic as a representation.

Manufacturing Intelligence

Wright, Paul Kenneth; and Bourne, David Alan
[New York University and
Carnegie Mellon University]

Reading, MA: Addison-Wesley, 1988, ix+352 pp
Hardbound, ISBN 0-201-13576-0

Reviewed by
Martin D. Levine
McGill University

If you are interested in the control of machine tools and how to apply artificial intelligence approaches to this problem you will find this book both interesting and useful. If you are interested in robotics, vision, or expert systems this book will only be worth consulting if you are involved in applications.

What do the authors mean by "manufacturing intelligence"? As indicated in the preface: 1) The goal of the emerging field of "manufacturing intelligence" is to model the skills and expertise of manufacturing craftsmen so that intelligent machines can make small batches of parts without human intervention. We believe that this goal will be accomplished by integrating the results of research in knowledge engineering, manufacturing software systems, robotic vision, and robotic manipulation. 2) It is the intention of this book to assess actual progress, in research and practice, towards full development and implementation of manufacturing intelligence in the factory. Further, we consider the rich promise that the book holds for bridging the two areas of manufacturing science and computer science.

The book contains lots of general discussions and comments regarding a wide range of aspects related to manufacturing intelligence. It also raises and discusses many important issues related to applying current research to practical industrial environments.

The book consists of four major parts. Part 1 has two chapters and deals with the machine tool industry. The authors set the tone for the book by specifying the desirability of automating craftsmanship, in particular that related to the operation of machine tools. This will clearly require an AI approach with particular emphasis on knowledge-based systems and machine perception of the environment. A discussion of manufacturing systems is presented and some desirable features for future systems are indicated.

Part 2 is a brief overview of intelligent machines for manufacturing from the point of view of design and contains four chapters. Chapter 3 is a general overview. Chapter 4 is concerned with the "manufacturing brain" and discusses software tools for building intelligent control systems. Software design is discussed in detail and the authors' Cell Management Language (CML) for programming manufacturing systems is presented. Chapter 5 provides a very restricted overview of the

"manufacturing eye". However, many pertinent questions related to industrial implementation of computer vision are raised and discussed from a design point-of-view. Chapter 6 is an interesting chapter on the evolution of robots with particular emphasis on the "manufacturing hand". The authors indicate the desirability of developing dextrous hands for grasping, and discuss both design guides and specific case studies.

Part 3 contains five chapters and deals specifically with the issues related to the development of an expert system for modeling "the skills of the craftsman in machining". Chapter 7 discusses a methodology for extracting the skills from the craftsman and would be of interest to people contemplating the design of expert systems for engineering applications. Chapter 8 is concerned with the design of an expert system for automated planning for machining. Chapter 9 discusses the new field of automated clamping and fixturing and presents several case studies. Chapter 10 is concerned with cutting sensors for machine tools, exemplified by a closed-loop control system which uses temperature and stress monitoring. Again case studies are presented and the authors indicate the large gap between what has been studied in research laboratories and what one can find on the factory floor. Chapter 11 discusses the team of experts required to develop intelligent machine tool control.

Part 4 is concerned with autonomous manufacturing and makes predictions about the state of manufacturing thirty years from now. It also includes a useful glossary of terms.

This book is not very useful as a general textbook because it is too focused on machine tools. On the other hand, if you are interested in a reference to and discussion of machine tools of the future, this book is definitely for you. It exhibits a real concern for applying advanced research to industrial applications and as such might interest "theory types" who would like to find out what the "real world" is concerned about.

Martin Levine is with the McGill Research Centre for Intelligent Machines, and is a Fellow of the Canadian Institute for Advanced Research.

Books Received

Books listed below that are marked ‡ will be reviewed in a future issue. Readers who wish to review books for the magazine should write, outlining their qualifications, to the book review editor, Graeme Hirst, Department of Computer Science, University of Toronto, Toronto, Ontario, Canada M5S 1A4. Obviously, we cannot promise the availability of books in anyone's exact area of interest.

Authors and publishers who wish their books to be considered for review in *Canadian AI* should send a copy to the book review editor at the address above. All books received will be listed, but not all can be reviewed.

Common Lisp: A Tutorial
Milner, Wendy L.
Englewood Cliffs: Prentice-Hall, 1988, xx+521 pp
Paperback, ISBN 0-13-152844-0

Computer Interpretation of Natural Language Descriptions
Mellish, Christopher S.
[University of Edinburgh]
(Ellis Horwood series in artificial intelligence)
Chichester, Eng.: Ellis Horwood Limited, 1985, 182 pp
Distributed by John Willey and Sons
Hardbound, ISBN 0-470-20219-X and 0-85312-828-6
US$26.95

Fundamentals in Computer Understanding: Speech and Vision: An Advanced Course
Haton, Jean-Paul (editor)
[Centre de Recherche en Informatique de Nancy]
Cambridge, Eng.: Cambridge University Press, 1987
xii+276 pp, Hardbound, ISBN 0-521-30983-2, US$39.50

Logic Programming and Knowledge Engineering
Amble, Tore
[University of Trondheim]
(International computer science series)
Wokingham, Eng.: Addison-Wesley, 1987, xii+281 pp
Paperback, ISBN 0-201-18043-X

Metataxis: Contrastive Dependency Syntax for Machine Translation
Schubert, Klaus
[BSO/Research, Utrecht]
(Distributed language translation 2)
Dordrecht and Providence, RI: Foris Pubs., 1987, 250 pp
Hardbound, ISBN 90-6765-358-6, US$37.00
Paperback, ISBN 90-6765-359-4, US$24.90

Natural Language Parsing and Linguistic Theories
Reyle, Uwe; and Rohrer, Christian (editors)
[University of Stuttgart]
(Studies in linguistics and philosophy 35)
Dordrecht: Reidel, 1988, vi+482 pp
Hardbound, ISBN 1-55608-055-7, US$79.00

Natural Language Parsing Systems
Bolc, Leonard (editor)
[Polish Academy of Sciences]
(Symbolic computation and artificial intelligence series)
Berlin: Springer-Verlag, 1987, xviii+367 pp
Hardbound, ISBN 3-540-17537-7 and 0-387-17537-7

Programming in Scheme
Eisenberg, Michael
[MIT Laboratory for Computer Science]
Redwood City: The Scientific Press, 1988, xv+304 pp
Paperback, ISBN 0-89426-115-0

‡ **A Robot Ping-Pong Player: Experiment in Real-time Intelligent Control**
Andersson, Russell L.
[AT&T Bell Laboratories]
(The MIT Press series in artificial intelligence)
Cambridge, MA: The MIT Press, 1988, xiii+275 pp
Hardbound, ISBN 0-262-01101-8, US$35.00

Self-organization and Associative Memory (second edition)
Kohonen, Teuvo
[Helsinki University of Technology]

(Springer series in information sciences 8)
Berlin: Springer-Verlag, 1988, xv+312 pp, Paperback
ISBN 3-540-18314-0 and 0-387-18314-0, US$39.50

‡ **What Every Engineer Should Know About Artificial Intelligence**
Taylor, William A.
Cambridge, MA: The MIT Press, 1988, xi+331 pp
Hardbound, ISBN 0-262-20069-4, US$25.00

Computational Intelligence

Abstracts for 4(1) February 1988
"Taking Issue"
Guest Editor: Mary McLeish

Editor's Note: These abstracts were not available at press time for the April issue.

LEARNING AND DISCOVERY: One System's Search for Mathematical Knowledge
S.L. Epstein
Hunter College
City University of New York
New York, New York

The Graph Theorist, GT, is a system which performs mathematical research in graph theory. From the definitions in its input knowledge base, GT constructs examples of mathematical concepts, conjectures and proves mathematical theorems about concepts, and discovers new concepts. Discovery is driven both by examples and by definitional form. The discovery processes construct a semantic net which links all of GT's concepts together.

Each definition is an algebraic expression whose semantic interpretation is a stylized algorithm to generate a class of graphs correctly and completely. From a knowledge base of these concept definitions, GT is able to conjecture and prove such theorems as "The set of acyclic, connected graphs is precisely the set of trees," and "There is no odd-regular graph on an odd number of vertices". GT explores new concepts either to develop an area of knowledge or to link a newly-acquired concept into a pre-existing knowledge base. New concepts arise from the specialization of an existing concept, the generalization of an existing concept, and the merger of two or more existing concepts. From an initial knowledge base containing only the definition of "graph", GT discovers such concepts as acyclic graphs, connected graphs and bipartite graphs.

The Consistency of Syntactical Treatments of Knowledge or **How to Compile Quantificational Modal Logics into Classical FOL**
Jim des Rivieres and Hector J. Levesque
Department of Computer Science
University of Toronto
Toronto, Ontario

The relative expressive power of a sentential operator a is compared to that of a syntactical predicate L(a) in the setting of first-order logics. Despite well-known results by Montague and by Thomason that claim otherwise, any of the so-called "modal" logics of knowledge and belief can be compiled into classical first-order logics that have a corresponding predicate on sentences. Moreover, through the use of a partial truth predicate, the standard modal axiom schemata can be translated into single sentences, making it possible to use conventional first-order logic theorem provers to directly derive results in a wide class of modal logics.

Logical Omniscience, Semantics and Models of Belief
Robert F. Hadley
School of Computing Science
Simon Fraser University
Burnaby, British Columbia

Logical omniscience may be described (roughly) as the state of affairs in which an agent explicitly believes anything which is logically entailed by that agent's beliefs. It is widely agreed that humans are not logically omniscient, and that an adequate formal model of belief, coupled with correct semantic theory, would not entail logical omniscience. Recently, two prominent models of belief have emerged which purport both to avoid logical omniscience and to provide an intuitively appealing semantics. The first of these models is due to Levesque (1984b); the second to Halpern (1985). It is argued herein that each of these models faces serious difficulties. Detailed criticisms are presented for each model and a computationally-oriented theory of intentions is presented which provides the foundation for a new formal model of belief. This formal model is presented in a decidable subset of first-order logic, and is shown to provide a solution to the general problem of logical omniscience. The model provides for the possibility of belief revision, and places no *a priori* restrictions upon an agent's representation language.

Abstracts for 4(3) August 1988
Special Issue on Planning
Guest Editor: David E.Wilkins

Hierarchical Planning Involving Deadlines, Travel Time and Resources
Thomas Dean
Department of Computer Science
Brown University
R. James Firby
Department of Computer Science
Yale University
David Miller
Department of Computer Science
Virginia Polytechnic Institute & State University

This paper describes a planning architecture that supports a form of hierarchical planning well suited to applications involving deadlines, travel time, and resource considerations. The architecture is based upon a temporal database, a heuristic evaluator, and a decision procedure for refining partial plans. A partial plan consists of a set of tasks and constraints on their order, duration, and potential resource requirements. The temporal database records the partial plan that the planner is currently working on, and computes certain

consequences of that information to be used in proposing methods to further refine the plan. The heuristic evaluator examines the space of linearized extensions of a given partial plan in order to reject plans that fail to satisfy basic requirements (e.g., hard deadlines and resource limitations), and estimates the utility of plans that meet these requirements. The information provided by the temporal database and the heuristic evaluator is combined using a decision procedure that determines how best to refine the current partial plan. Neither the temporal database nor the heuristic evaluator is complete, and, without reasonably accurate information concerning the possible resource requirements of the tasks in a partial plan, there is a significant risk of missing solutions. A specification language that serves to encode expectations concerning the duration and resource requirements of tasks greatly reduces this risk, enabling useful evaluations of partial plans. Details of the specification language and examples illustrating how such expectations are exploited in decision making are provided.

Exploiting Temporal Coherence in Nonlinear Plan Construction
Mark Drummond and Ken Currie
AI Applications Institute
University of Edinburgh
Edinburgh, U.K.

Correct conventional nonlinear planners operate in accordance with Chapman's (1985a) modal truth criterion (MTC). The MTC characterizes the conditions under which an assertion will be true at a point in a nonlinear plan. However, the MTC is not all one requires in order to build a realistic planning system: it merely sanctions the use of a number of plan modifications in order to achieve each assertion in a developing plan. The number of modifications that can be made is usually very large. To avoid breadth first search a planner must have some idea of which plan modification to consider. We describe a domain independent search called temporal coherence which helps guide the search through the space of partial plans defined by the MTC. Temporal coherence works by suggesting certain orderings of goal achievement as more appealing than others, and thus by finding bindings for plan variables consistent with the planner's overall goals. Our experience with a real nonlinear planner has highlighted the need for such a heuristic. In this paper, we give an example planning problem, and use it to illustrate how temporal coherence can speed the search for an acceptable plan. We also prove that if a solution exists in the partial plan search space defined by the MTC then there exists a path to that solution which is sanctioned by temporal coherence.

Localized Event-Based Reasoning for Multiagent Domains
Amy L. Lansky
Artificial Intelligence Center
SRI International
Menlo Park, CA
and
Center for Study of Language and Information
Stanford University
Stanford, CA

This paper presents the GEM concurrency model and GEMPLAN, a multiagent planner based on this model. Unlike standard state-based AI representations, GEM is unique in its explicit emphasis on events and domain structure. In particular, a world domain is modeled as a "set of regions composed of interrelated events". Event-based temporal logic constraints are then associated with each region to delimit legal domain behaviour. The GEMPLAN planner directly reflects this emphasis on domain structure and constraints. It can be viewed as a general purpose constraint satisfaction facility which constructs a network of interrelated events (a "plan") that is subdivided into regions ("subplans"), satisfies all applicable regional constraints, and also achieves some stated goal.

GEMPLAN extends and generalizes previous planning architectures in the range of constraint forms it handles and in the flexibility of its constraint satisfaction search strategy. One critical aspect of our work has been an emphasis on localized reasoning — techniques that make explicit use of domain structure. For example, GEM localizes its applicability of domain constraints and imposes additional "locality constraints" on the basis of domain structure. Together, constraint localization and locality constraints provide semantic information that can be used to alleviate several aspects of the frame problem for multiagent domains. The GEMPLAN planner reflects the use of locality by subdividing its constraint satisfaction search space into "regional planning search spaces". Utilizing constraint and property localization, GEMPLAN can pinpoint and rectify interactions among these regional search spaces, thus reducing the burden of "interaction analysis" ubiquitous to most planning systems. Because GEMPLAN is specifically geared toward parallel, multiagent domains, we believe that its natural application areas will include scheduling and other forms of organizational coordination.

Causal Reasoning in Planning
David E. Wilkins
Artificial Intelligence Center
SRI International
Menlo Park, CA

Reasoning about actions necessarily involves the truth of assertions about the world over time. The SIPE planning system retains the efficiency of the STRIPS assumption for this while enhancing expressive power by allowing the specification of a causal theory. Separation of knowledge about causality from knowledge about actions relieves operators of much of their representational burden and allows them to be applicable in a wide range of contexts. The implementation of causal theories is described, together with examples and evaluations of the system's power and efficiency.

Synthesizing Plans that Contain Actions with Context-Dependent Effects
Edwin P.D. Pednault
Knowledge Systems Research Department
AT&T Bell Laboratories
Holmdel, New Jersey

This paper presents a method of solving planning problems that involve actions whose effects change according to the situations in which they are performed. The approach is an extension of the conventional planning methodology in which plans are constructed

through an iterative process of scanning for goals that are not yet satisfied, inserting actions that produce the same effects in every situation. The extension involves introducing additional subgoals to actions above and beyond the preconditions of execution normally introduced. These additional subgoals, called secondary preconditions, ensure that the actions are performed in contexts conducive to producing the effects we desire. This paper defines and analyzes secondary preconditions from a mathematically rigorous standpoint and demonstrates how they can be derived from regression operators.

Subscription Information

Computational Intelligence is published by the National Research Council and is sponsored by CSCSI/SCEIO. Non-institutional CSCSI/SCEIO members may subscribe for Cdn$16.00, half the regular price, by using a copy of the order form on the last page of this magazine.

Technical Reports

University of Saskatchewan

Intelligent Advising in Problem Solving Domains: The SCENT-3 Architecture
Gordon I. McCalla and Jim E. Greer
ARIES Laboratory Research Report # 88 - 1
(Laboratory for Advanced Research in Intelligent Educational Systems)
Dept. of Computational Science
University of Saskatchewan
Saskatoon, Canada

The SCENT project has focused on developing an intelligent advising environment for students learning to program in Lisp. This paper presents a status report on SCENT, a summary of recent research contributions, and the design of the new SCENT-3 system. SCENT-3 is an architecture for a full scale student advising system. Work on SCENT-3 is progressing concurrently on many fronts with new research contributions in dynamic planning and blackboard control, student modeling, strategy judging and diagnosis, and program analysis. The SCENT-3 architecture has been designed to achieve the goal of creating an intelligent advising system for higher order problem-solving activities with minimal domain dependence. Much more research remains to be done to fully prove out SCENT-3, but the investigations carried out so far are promising.

Student Models: The Genetic Graph Approach
Barbara Brecht and Marlene Jones
ARIES Laboratory Research Report # 88 - 2
(Laboratory for Advanced Research in Intelligent Educational Systems)
Dept. of Computational Science
University of Saskatchewan
Saskatoon, Canada

In this paper we examine the student model component of an Intelligent Computer-Assisted Instruction (ICAI) system. First, we briefly discuss the desirable capabilities of the student model and then describe, in detail, one approach to student modeling which is based on Goldstein's genetic graph. We expand Goldstein's definition and test it's feasibility in new domains, since his original domain was a limited, straightforward adventure game. In addition to modeling two diverse domains, subtraction and ballet, we also discuss the role of certain ICAI components in generating and maintaining the genetic graph.

Problem Solving by Analogy in Novice Programming
Judy A. Escott
ARIES Laboratory Research Report # 88 - 3
(Laboratory for Advanced Research in Intelligent Educational Systems)
Dept. of Computational Science
University of Saskatchewan
Saskatoon, Canada

Problem solving by analogy is a domain independent cognitive strategy that can provide a reference for solving new problems that are similar to previously solved problems and can encourage generalization of domain specific knowledge. Novices are prime candidates to make incorrect analogies because of limited domain knowledge and limited abstract understanding of that knowledge. The goals of this thesis are to demonstrate that incorrect analogy is a source of error in novice programming, to discern the advantages of having an intelligent tutoring system identify incorrect analogies, and to illustrate how incorrect analogy detection can be operationalized by using the student history, the current solution analysis and knowledge of the types of incorrect analogies students make. Results of the analysis of novices' solutions to four recursive Lisp problems provide evidence to support incorrect analogy as a cause of errors and identify the types of incorrect analogies made: transfer of specific strategies, transfer of and superficial modification of strategies to suit different problem requirements, and transfer of specific strategy components to implement different strategies. An operational approach to detecting incorrect analogies is a first step toward tutoring this type of error.

Editors' Note: This is Judy Escott's M.Sc. thesis, for which she has won the Governor General's Gold Medal for the most distinguished Master's student graduating in the Spring, 1988 Convocation at U of S.

Ordering information

Requests for any of the aforenamed publications should be addressed to: Jim Greer, ARIES Laboratory, Dept. of Computational Science, University of Saskatchewan, Saskatoon, Saskatchewan, Canada, S7N 0W0.

Deadline for the October issue is 15 August

CALL FOR PAPERS

FIRST INTERNATIONAL CONFERENCE ON PRINCIPLES OF KNOWLEDGE REPRESENTATION AND REASONING

Royal York Hotel
Toronto, Ontario, Canada

May 15-18, 1989

Sponsored by the Canadian Society for Computational Studies of Intelligence,
with support from AAAI, IJCAI, the Canadian Institute for Advanced Research, and the Information Technology Research Centre of Ontario,
in cooperation with AISB and ACM SIGART (pending approval)

The idea of explicit representations of knowledge, manipulated by general-purpose inference algorithms, underlies much of the work in artificial intelligence, from natural language to expert systems. A growing number of researchers are interested in the principles governing systems based on this idea. This conference will bring together these researchers in a more intimate setting than that of the general AI conferences. Authors will be expected to give presentations of adequate length to present substantial results, and parallel sessions will be avoided to the extent possible. Accepted papers will be collected in a conference proceedings, to be published by Morgan Kaufmann Publishers, Inc.

The conference will focus on principles of commonsense reasoning and representation, as distinct from concerns of engineering and details of implementation. Thus of direct interest are logical specifications of reasoning behaviors, comparative analyses of competing algorithms and theories, and analyses of the correctness and/or the computational complexity of reasoning algorithms. Papers that attempt to move away from or refute the knowledge-based paradigm in a principled way are also welcome, so long as appropriate connections are made to the central body of work in the field.

Submissions are encouraged in at least the following topic areas:

- Analogical Reasoning
- Commonsense Reasoning
- Deductive Reasoning
- Diagnostic and Abductive Reasoning
- Evidential Reasoning
- Inductive Reasoning
- Nonmonotonic Reasoning
- Qualitative Reasoning
- Temporal Reasoning
- Planning
- Knowledge Representation Formalisms
- Theories of the Commonsense World
- Theories of Knowledge and Belief
- Belief Management and Revision
- Formal Task and Domain Specifications

REVIEW OF PAPERS

The Program Committee will review *extended abstracts* (not complete papers). Submissions will be judged on clarity, significance, and originality. An important criterion for acceptance is that the paper clearly contribute to principles of representation and reasoning that are likely to influence current and future AI practice.

Extended abstracts should contain enough information to enable the Program Committee to identify the principal contribution of the research and its importance. It should also be clear from the extended abstract how the work compares to related work in the field. References to relevant literature must be included.

Submitted papers must be unpublished. Submissions must also be substantively different from papers currently under review and must not be submitted elsewhere before the author notification date (December 15, 1988).

SUBMISSION OF PAPERS

Submitted abstracts must be at most *eight (8) double-spaced pages.* All abstracts must be submitted on 8-1/2" x 11" paper (or alternatively, a4), and printed or typed in *12-point font* (pica on standard typewriter). Dot matrix printout is not acceptable.

Each submission should include the *names* and *complete addresses* of all authors. Also, authors should indicate under the title which of the *topic areas* listed above best describes their paper (if none is appropriate, please give a set of keywords that best describe the topic of the paper).

Abstracts must be *received* no later than November 1, 1988, at the address listed below. Authors will be notified of the Program Committee's decision by December 15, 1988.

Authors of accepted papers will be expected to submit substantially longer full papers for the conference proceedings. Final camera-ready copies of the full papers will be due on February 15, 1989. Final papers will be allowed at most twelve (12) double-column pages in the conference proceedings.

Send five (5) copies of extended abstracts [one copy is acceptable from countries where access to copiers is limited] to

Ron Brachman and Hector Levesque, Program Co-chairs
First Int'l. Conference on Principles of Knowledge Representation and Reasoning
c/o AT&T Bell Laboratories
600 Mountain Avenue, Room 3C-439
Murray Hill, NJ 07974
USA

INQUIRIES

Inquiries of a general nature can be addressed to the Conference Chair:

Raymond Reiter, Conference Chair
First Int'l. Conference on Principles of Knowledge Representation and Reasoning
c/o Department of Computer Science
University of Toronto
10 Kings College Road
Toronto, Ontario M5S 1A4
CANADA
electronic mail: reiter@ai.toronto.edu

IMPORTANT DATES

Submission receipt deadline:	November 1, 1988
Author notification date:	December 15, 1988
Camera-ready copy due to publisher:	February 15, 1989
Conference:	May 15–18, 1989

PROGRAM COMMITTEE

James Allen (University of Rochester)
Giuseppe Attardi (Delphi SpA, Italy)
Woody Bledsoe (MCC/University of Texas)
Alan Bundy (Edinburgh University)
Eugene Charniak (Brown University)
Veronica Dahl (Simon Fraser University)
Koichi Furukawa (ICOT)
Johan de Kleer (Xerox PARC)
Herve Gallaire (ECRC, Munich)
Michael Genesereth (Stanford University)
Michael Georgeff (SRI International)
Pat Hayes (Xerox PARC)
Geoff Hinton (University of Toronto)
Bob Kowalski (Imperial College)
Vladimir Lifschitz (Stanford University)
Alan Mackworth (University of British Columbia)
Drew McDermott (Yale University)
Tom Mitchell (Carnegie-Mellon University)
Robert Moore (SRI International)
Judea Pearl (UCLA)
Stan Rosenschein (SRI International)
Stuart Shapiro (SUNY at Buffalo)
Yoav Shoham (Stanford University)
William Woods (Applied Expert Systems)

Upcoming Conferences

In Canada

First International Conference on Principles of Knowledge Representation and Reasoning
15 - 18 May 1989, Toronto, Ontario
See announcement on pages 42 - 43 for details.

3rd AAAI Workshop on Knowledge Acquisition for Knowledge-Based Systems
7 - 11 November 1988, Banff, Alberta
Topics include: Transfer/modeling of expertise; Learning systems; Extracting and modeling of knowledge; Integration of Knowledge Acquisition (KA) techniques and systems; KA methodology and training; Validation of KA techniques.
Contact: John Boose, Advanced Technology Center, Boeing Computer Services, 7L-64; via mail: PO Box 24346, Seattle, Washington, USA 98124; via courier: Bldg 33.07 2760 160th Ave. SE, Bellevue, Washington, USA 98008. Phone: (206) 865-3253. Brian Gaines, Dept. of Comp. Sci., U. of Calgary, 2500 University Dr. NW, Calgary, Alberta T2N 1N4. Phone: (403) 220-5901.

6th Canadian Symposium on Instructional Technology
3 - 5 May 1989, Halifax, Nova Scotia
Focus: Computer-assisted learning — theory and reality.
Topics include: AI in education training; Innovations in instructional technology and courseware development; Technology transfer from researcher to user; Impact of ES and AI on CAL; User perspective on CAL, implementation and application.
Submission material: Abstract (half page absolute max.).
Submission deadline: Sept 30, 1988.
Contact: F. Kewley, Sixth Canadian Symposium on Instructional Technology, Conference Services Office, NRC Canada, Ottawa, Ontario, K1A 0R6. Phone:(613) 993-9009. Telex: 053-3145.

In the United States

IEEE ICNN-88: Conference on Neural Networks
23 - 27 July 1988, San Diego, California
Sessions include: Self-organization; Network architectures; Network dynamics; Learning algorithms; Associative memory; Image processing applications; Vision; Optical neurocomputers; Neurobiological connections; Combinatorial optimization; Speech recognition and synthesis; Robotics.
Contact: Nomi Feldman, IEEE ICNN-88 Conference Secretariat, 3770 Tansy St., San Diego, CA 92121.
Phone: (619) 453-6222.

5th ACM Conference on Lisp and Functional Programming
25 - 27 July 1988, Snowbird, Utah
Topics include: Programming language concepts and facilities; Implementation methods; Machine architectures; Semantic foundations; Programming logics; Program development environments.
Contact: Robert Cartwright, Dept of Comp. Sci., Rice U., P. O. Box 1892, Houston, TX 77251-1892. Phone: (713) 527-4834. NET: cork@rice.edu.

Workshop on Computational Learning Theory
3 - 5 August 1988, Cambridge, Massachusetts
Topics include: Resource, convergence-rate and robustness analysis of specific learning algorithms; General learnability and non-learnability results in existing models and general upper and lower bounds on resources required for learning; New models, extensions of existing models, and theoretical comparisons among models. Papers that make formal connections with work in robotics, neural nets, pattern recognition, adaptive signal processing and cryptography are also welcome. Financial support may be available for graduate students.
Contact: John Cherniavsky, Workshop on Computational Learning Theory, Dept of Comp. Sci., Georgetown U. Washington, D.C. 20057.

DIAC-88: Directions and Implications of Advanced Computing
21 August 1988, St. Paul, Minnesota
Topics include: Ethical issues in computing research; Sources and effects of research funding; Responsible software development; AI and the conduct of war; Limits to the automation of war; Automated defense systems; Computerized voting; Civil liberties; Risks of the new technology; Resource modeling; Arbitration and conflict resolution; Software safety.
Contact: Nancy Leveson, ICS Department, U. of California, Irvine, CA 92717. Phone: (714) 856-5517. Sponsored by Computer Professionals for Social Responsibility, P.O. Box 717, Palo Alto, CA 94301.

AAAI-88: 7th National Conference on AI
22 - 26 August 1988, St. Paul, Minnesota
Topics include: AI and education; Automated reasoning; Cognitive modeling; Commonsense reasoning; Expert systems; Knowledge acquisition and representation; Machine architecture and computer languages for AI; Machine learning; Natural language; Robotics; User interfaces.
Contact: AAAI-88, American Association for Artificial Intelligence, 445 Burgess Dr., Menlo Park, CA 94025-3496.

AAAI-88 Workshops
19 - 27 August 1988, St. Paul, Minnesota
Uncertainty in AI: *19 - 21 Aug*
Contact: Ross Shachter, Center for Health Policy, 125 Old Chemistry Bldg., Duke U., Durham NC 27706.
Phone: (919) 684-4424.

Email: shachter@sumex-aim.stanford.edu.

Neural Architectures for Computer Vision: *20 Aug*
Contact: Patrick Ransil, Lockheed AI Center, 2710 Sand Hill Rd., Menlo Park, CA 94025.

Parallel Algorithms for Machine Intelligence and Pattern Recognition: *20 - 21 Aug*
Contact: Prof. Laveen N. Kanal, Dept. of Comp. Sci., U. of Maryland, College Park, MD 20742.
Email: kanal@mimsy.umd.edu.

Integration of Knowledge Acquisition and Performance Systems: *21 Aug*
Contact: Brian R. Gaines, Dept. of Comp. Sci., U. of Calgary, Calgary, Alberta T2N 1N4. Phone: (403) 220-5901. Email: gaines@calgary.cdn.

Explanation: *22 Aug*
Contact: Michael R. Wick, Comp. Sci. Dept., U. of Minnesota, Minneapolis, MN 55455.

AI and Hypertext: Issues and Directions: *23 Aug*
Contact: Mark Bernstein, Eastgate Systems Inc., PO Box 1307, Cambridge, MA 02238. Phone: (617) 782-9044.

Case-Based Reasoning: *23 Aug*
Contact: Edwina L. Rissland, Dept. of Comp. and Info. Sci., U. of Massachusetts, Amherst, MA 01003.

AI in Design: *24 Aug*
Contact: D. Navinchandra, Robotics Institute, Carnegie-Mellon U., Pittsburgh, PA 15213. Phone: (412) 268-8825. Email: dchandra@isl1.ri.cmu.edu.

AI and Music: *24 Aug*
Contact: Kemal Ebcioglu, H2-L10, IBM, Thomas J. Watson Research Center, PO Box 704, Yorktown Heights, NY 10598.
Email: kemal@ibm.com., kemal@ytkvmh2.bitnet.

Blackboard Systems: *24 Aug*
Contact: V. Jagannathan, M/S 7L-64, Boeing Advanced Technology Center, Boeing Computer Services, PO Box 24346, Seattle, WA 98124-0346. Phone: (206) 865-3240. Email: juggy@boeing.com.

AI in Process Engineering: *25 Aug*
Contact: Michael Mavrovouniotis, Massachusetts Institute of Technology, Room 66-056, 77 Massachusetts Ave., Cambridge, MA 02139.

AI in Production Planning and Scheduling: *25 Aug*
Contact: Mark S. Fox, Robotics Institute, Carnegie-Mellon U., Pittsburgh, PA 15213. Phone: (412) 268-3832. Email: msf@isl1.ri.cmu.edu.

Automating Software Design: Current Directions: *25 Aug*
Contact: Douglas R. Smith, Kestrel Institute, 1801 Page Mill Rd., Palo Alto, CA 94304-1216. Phone: (415) 493-6871. Email: smith@kestrel.arpa.

Databases in Large AI Systems: *26 Aug*
Contact: Forouzan Golshani, Dept. of Comp. Sci., Arizona State U., Tempe, AZ 85287. Phone: (602) 965-2855.

Conceptual Graphs: *27 Aug*
Contact: Dr. John W. Esch, Unisys Defense Systems, PO Box 64525 MS CCA06, St. Paul, MN 55164-0525.

3rd IEEE International Symposium on Intelligent Control

24 - 26 August 1988, Arlington, Virginia

Topics include: Adaptive and self-organizing controllers; Hierarchical control systems; Sensor-based robot control; Knowledge-based control systems; Qualitative reasoning in intelligent control; Models of approximate reasoning; Fuzzy control systems; CAD; Linguistic control systems; Multiple agent systems; cooperative and antagonistic systems; Neural networks as a control tool; Expert databases for intelligent control; Communication protocols for intelligent manufacturing systems; Space and underwater exploration.
Contact: Dr. Hany K. Eldeib, Dept of Electrical and Comp. Engineering, George Mason U., 4400 University Dr., Fairfax, VA 22030. Phone: (703) 323-3369. Bitnet: heldeib@gmuvax. UUCP: uunet!pyrdc!gmu90x!heldeib.

International Neural Network Society Annual Meeting

6 - 10 September 1988, Boston, Massachusetts

Contact: Neural Networks, AT&T Bell Labs, Room 4G-323, Holmdel, NJ 07733. Second Annual meeting will be held September 5 - 9, 1989, Washington, D.C.

4th Annual Expert Systems in Government Conference

October 1988, Washington, D.C.

Theme: Intelligent Systems — realizing the payoff for today & tomorrow. Topics include: Knowledge representation and acquisition; Reasoning under uncertainty; Natural language and intelligent interfaces; Expert system development environments.
Contact: ESIG '88, MS W418, The Mitre Corporation, 7525 Colshire Dr., McLean, VA 22102.

4th Aerospace Applications of AI Conference

25 - 27 October 1988, Dayton, Ohio

Tutorials will be held Oct 24 and workshops on Oct 28.
Topics include: Integrating neural networks and Expert Systems (ES); Machine learning, cognition and the cockpit; Neural networks and human-machine interfaces; Parallel processing and Neural Networks; Back propagation with momentum, shared weights or recurrent; ES development tools; Aerospace scheduling; Real-time expert systems; Verification and validation of ES; Natural language recognition and synthesis.
Contact: James Johnson, AFWAL/AAOR, WPAFB, OH 45433.

2nd IEEE Conference on Neural Information Processing Systems (Natural and Synthetic)

28 November - 1 December 1988, Denver, Colorado

Topics include: Neurobiological models of development; Cellular information processing; Synaptic function, learning and memory; Connectionist models of learning and cognitive processing; Training paradigms; Analysis of applicability; Generalization and complexity; Applications to signal processing, vision, speech, motor control, knowledge engineering and adaptive systems; Advances in hardware technologies — neurophysiological recording tools, VLSI or optical implementations of neural networks.
Contact: Scott Kirkpatrick, IBM T. J. Watson Research Center, P.O. Box 704, Yorktown Heights, NY 10598.

2nd Conference on Computer Vision

5 - 8 December 1988, Tarpon Springs, Florida

Focus: All aspects of computer vision.
Contact: ICCV'88, c/o Computer Society of the IEEE, 1730 Massachusetts Ave., N.W., Washington, D.C. 20036-1903.

5th Conference on Data Engineering

7 - 9 February 1989, Los Angeles, CA

Topics include: AI and knowledge based systems;

Applications and application systems; Autonomous distributed systems; Concurrency control and data integrity; Data access control and security; Data engineering techniques and tools; Data services and servers; Performance evaluation; Communication systems.
Contact: John Carlis, Comp. Sci. Dept., U. of Minnesota, 207 Church St, SE, Minneapolis, MN 55455. Phone: (612) 625-6092. NET: carlis%umn-cs.arpa@relay.cs.net.

ICCAL '89: 2nd Conference on Computer-Assisted Learning

9 - 11 May 1989, Dallas, Texas

Topics include: AI applications and instruction; Intelligent tutoring systems; Knowledge acquisition and representation; Student modeling and cognitive diagnosis; Human computer instruction; Computational models of reasoning and learning; Evaluation of learning environments; Knowledge based CAI systems; Authoring systems; Performance monitoring.
Submission material: 4 copies of extended abstract (1500 - 5000 words). Submission deadline: Sept 15, 1988.
Contact: Prof. Hermann Maurer, IIG, Schiesstattgasse 4a, A-8010 Graz, Austria. Phone: 0043-316-70255/12.
Email: maurer@btx.uucp.

IJCAI 11th Joint Conference on AI

20 - 25 August 1989, Detroit, Michigan

The technical program consists of a Paper Track focusing on empirical, analytical, theoretical, conceptual, foundational aspects and applied research. The Videotape Track focuses on applications in all subfields best suited for this type of presentation.
Topics include: AI tools and technologies; Machine architectures, languages, shells; Search methods; Knowledge acquisition, learning, analogy; Real-time performance; Parallel and distributed processing; Cognitive modeling; Planning, scheduling, reasoning about actions; Natural language, speech understanding and generation; Perception, vision, robotics; Intelligent tutoring systems; Design, manufacturing, control; Philosophical foundations, perspectives and attitudes, social implications.
Submission deadline: Dec 7, 1988. Submit 6 copies in hardcopy, of 4-10 single spaced pages, with 100-200 word abstract. Submit 1 copy of 15 minute video, with written abstract and indication of tape format (NTSC, PAL, SECAM; VHS, .75"U-matic).
Contact: IJCAI 89, c/o AAAI, 445 Burgess Drive, Menlo Park, CA 94025-3496.

Outside North America

2nd Workshop on Qualitative Physics

26 - 28 July 1988, Paris, France

Focus: A forum for discussion of ongoing research in qualitative physics and related areas. Topics include: Causal reasoning; Mathematical aspects of qualitative models; Naive physics versus qualitative physics. Attendance by invitation only. Deadline to apply was March 8.
Contact: Francesco Gardin, Dipartimento di Scienze dell'Informazione, Universita degli Studi di Milano, Via Moretto da Bresica, 9 20133 Milano, ITALY.
Phone: +39-2-2141230.

4th Conference on Modeling Techniques and Tools for Computer Performance Evaluation

15 - 17 September 1988, Mallorca, Spain

Topics include: Exploration of Expert System techniques in system modeling, configuration and monitoring; Experiences in workload characterization for databases, distributed systems and AI applications.
Contact: D. Potier, Departament de Ciencies Matematiques i Informatica, Universitat de les Illes Balears, Miquel dels Sants Oliver 2, 07012 Palma de Mallorca, Mallorca, Spain. Phone: (34) (71) 292440.

EWSL 88: 3rd European Working Session on Learning

3 - 5 October 1988, Glasgow, Scotland

Focus: The emphasis will be on Machine Learning, but relevant Cognitive Science studies are welcomed.
Contact: Derek Sleeman, Dept of Comp. Sci., U. of Aberdeen., ABERDEEN AB9 2UB, Scotland UK. Phone: Aberdeen (+44 224) 272288. Telex: 73458.

1st European Congress on AI and Training

4 - 6 October 1988, Lille, France

Focus: To promote and develop contacts between various fields which participate in design and creation of AI systems for computerized teaching (AI, didactic, cognitive psychology).
Topics include: Dialogue and natural language; Cognitive process of understanding and knowledge acquisition; Learning, by people and machine; Reasoning; Expert system shells and languages; Simulation; Intelligent tutoring systems; Didactic interference, explanation; Design and estimation methodology.
Contact: Congres Applica, Secretariat Scientifique, C.C.I.L.R.T. B.P. 359 - 59020 Lille Cedex, France.

IAPR Workshop on Computer Vision Special Hardware and Industrial Applications

12 - 14 October 1988, Tokyo, Japan

(This workshop will be held in advance of)

9th Conference on Pattern Recognition (ICPR)

17 - 20 October 1988, Beijing, China

Topics include: Special hardware and industrial applications; High speed image processor; VLSI image signal processor chip; PC-based low-cost image analysis system; Special-purpose PRIP machine; Intelligent sensor; Visual inspection; Robot vision; Engineering automation for documents and line drawings; New imaging techniques; 3-D information usage.
Contact: Mikio Takagi, Institute of Industrial Science, U. of Tokyo, 7-22-1, Roppongi, Minato-ku, Tokyo 106.
CSNET: takagi%iis.u-tokyo.junet@relay.cs.net.

1st Symposium on AI

24 - 28 October 1988, Monterrey, Mexico

Topics include: Knowledge-based systems; Knowledge acquisition; Knowledge representation; Inference engine; Certainty factors; Vision; Robotics; Expert Systems applications in industry; Natural language processing; Speech recognition.
Submission material: 4 copies of summaries (5 pages max.). Submission deadline: August 31, 1988.
Contact: ITESM, Centro de Investigacion en Informatica, David Garza Salazar, Sucursal de Correos J. 64849 Monterrey, N.L. Mexico. Phone: (83) 59 57 47, (83) 59 59 43, (83) 59 57 50. BITNET: SIIACII AT TECMTYVM.

1st Australian Knowledge Engineering Congress
2 - 4 November 1988, Melbourne, Australia
Topics include: Expert systems case studies; Knowledge engineering methodologies; Design and use of conceptual schemas; Natural language interfaces; Evaluation of tools and expert systems; Role of consultants in knowledge engineering; Design of intelligent tutors.
Submission deadline: "ASAP". Submission material: A preliminary indication of interest in offering a paper.
Contact: B.J. Garner, Deakin U., Victoria 3217, Australia.

AI'88: Australian Joint AI Conference
15 - 18 November 1988, Adelaide, Australia
Focus: To bring together business, industry and researchers; To help business and industry understand the potential in AI; To provide a forum for research topics; To gain insight into the process of going from research to product; To be a joint, multi-theme conference.
Contact: AI'88 Secretariat, Dept of Comp. Science, U. of Adelaide, GPO Box 498, Adelaide, South Australia 5001, Australia. Phone: (08) 228-5586.
Telex: UNIVAD AA89141. Fax: (08) 224-0464. Internet: AI88@uacomsci.ua.oz.au.

Neural Networks and Applications
15 - 17 November 1988, Nanterre, France
Topics include: Languages, models, simulation, tools, implementation techniques; Vision and image processing; Acoustics; Speech recognition; Character recognition; Robotics; Learning; Knowledge aquisition; Optimization and combinatorics; Diagnostic techniques and quality control; Surveillance, security, reliability, fault tolerance.
Contact: Workshop Secretary, EC2, Neuro-Nimes, 269, rue de la Garenne, 92000 Nanterre, France. Phone: (1) 27 80 70 00. Telex: 612 469 F. Fax: (1) 47 80 66 29.

International Computer Science Conference '88
AI: Theory and Applications
19 - 21 December 1988, Hong Kong
Topics include: AI architectures; Expert systems; Knowledge engineering; Logic programming; Machine learning; Natural languages; Neural networks; Pattern recognition; Robotics; CAD/CAM; Chinese computing; Distributed systems; Information systems; Office automation; Software engineering.
Contact: Jean-Louis Lassez, Rm H1-A12, IBM Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights NY 10598. E-mail: JLL@ibm.com.

2nd Workshop on AI in
Economics and Management
11 - 13 January 1989, Singapore
Focus: Finance, banking, insurance, economics, decision support systems, public and private services, office automation, law, manufacturing planning, personnel and assets administration. Submission may be a paper or a demonstration of AI software program.
Submission deadline: July 1, 1988. Submission material: 2 copies of 700 word extended abstract.
Contact: Desai Narasimhalu, Institute of Systems Science, National U. of Singapore, Kent Ridge, Singapore 0511, Singapore. BITNET: issad@nusvm.

4th Conference on AI and Education
24 - 26 May 1989, Amsterdam, The Netherlands
Topics include: Intelligent tutoring systems (ITS); Development methods of ITS; Relevant cognitive and educational research; Advanced ITS Architectures; Domain Representation; Student modeling and diagnosis of student problems; Interaction/teaching strategies; Evaluation of ITS; AI-based learning environments; Modeling/simulating worlds; AI languages in the classroom.
Submission deadline: December 1988.
Contact: AIED '89 Secretary, SWI, U. of Amsterdam, Herengracht 196, 1016 BS Amsterdam, The Netherlands.
Email: aied@mcvax!swivax.uucp.

6th Scandinavian Conference on Image Analysis
19 - 22 June 1989, Oulu, Finland
Topics include: Computer vision; Perception; Image processing; Parallel algorithms; Pattern recognition; Applications.
Submission deadline: Dec 1,1988. Submission materials: 4 copies of extended summary (min. 1000 words).
Contact: Prof. Matti Pietikainen, 6SCIA Program Chairman, Dept. of Electrical Engineering, U. of Oulu, SF-90570 OULU, Finland. Phone: +358-81-352765. Fax: +358-81-561278. Telex: 32 375 oylin sf. Net: scia@steks.oulu.fi.

CSCSI/SCEIO Application and Order Forms

Use the forms below to subscribe to the journal *Computational Intelligence*, to order publications and to join the Canadian Society for Computational Studies of Intelligence (CSCSI/SCEIO) with which you will receive *Canadian Artificial Intelligence*. Complete the form of your choice and send it to CIPS (which administers membership for the CSCSI/SCEIO) at the address below, with the appropriate fee. (NOTE: Those residing outside of Canada who wish to order the 1988 CSCSI/SCEIO Conference Proceedings must mail the conference order form to Morgan Kaufmann Publishers.):

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

CSCSI/SCEIO Membership
Canadian Artificial Intelligence Back Issues

( ) I wish to join CSCSI/SCEIO and receive *Canadian Artificial Intelligence* (Cdn$35).
( ) I am a student (subtract Cdn$10). **AND/OR** ( ) I am a member of CIPS (subtract Cdn$10).

( ) Please send me the following back issues of *Canadian Artificial Intelligence* (Cdn$10 each including postage, except Cdn$15 for #3):*.....

Name .....

Mailing Address .....

.....

*If an issue you request is out of print, a photocopy will be provided. Issue #3 (March 1985) includes the supplement *Towards a Canadian Fifth Generation Research Plan.*

Computational Intelligence Subscriptions

( ) Please enter my subscription to *Computational Intelligence* at the CSCSI/SCEIO non-institutional member discount rate (Cdn$16).‡

Name .....

Mailing Address .....

.....

‡*Computational Intelligence* subscriptions are filled by its publisher, the National Research Council of Canada. CIPS only certifies your eligibility for the discount and forwards your order.

CSCSI/SCEIO Conference Proceedings

( ) Please send me the following CSCSI/SCEIO conference proceedings (indicate number of copies desired):

( ) Edmonton, 1988 (Within Canada: Members: Cdn$35; Non-members: Cdn$40. Add Cdn$5 for postage. Mail to **CIPS**.) (Outside Canada: US$35. Postage within U.S.: US$2 for the first copy, US$.75 for each additional copy. Postage outside U.S.: US$3 for the first copy, US$2 for each additional copy. Within California: add 7% sales tax. Mail to: **Morgan Kaufmann Publishers**, Order Fulfillment Center, PO Box 50490, Palo Alto, CA 94303, USA.)

( ) Montréal, 1986 (Cdn$30. Postage within Canada: Cdn$5. Outside Canada: Cdn$7. Mail to **CIPS**.)

( ) Saskatoon, 1982 ( ) Victoria, 1980 (Cdn$25. Postage within Canada: Cdn$5. Outside Canada: Cdn$7. Mail to **CIPS**.)

Name .....

Mailing Address .....

.....

Applied AI Systems, Inc.

A Window Into The Future

- Knowledge on the frontier of AI research
- Knowledge on real-time, real-world AI systems
  - comprehensive survey of the state of the art
  - international research project underway
- Knowledge on international AI business
  - over 1700 research and business contacts in Canada, the U.S., Britain, Europe, Japan and Asia
  - most major AI research institutions visited
  - contracts from major international AI R&D organizations
- Knowledge about world 6th Generation Computing
  - paper presented at AI-II Workshop, sponsored by Transport Canada, March 1988
  - paper presented at an international symposium on AI, in Europe
  - research contracts from Japanese 6th Generation research laboratories
- Knowledge about building AI systems for government applications
  - paper presented at Expert Systems in Government Conference, Washington, Dayton, Ottawa, Japan and Germany
  - several systems built or being built for government agencies

Applied AI Systems, Inc.
Gateway Business Park
300 March Road, Suite 602
KANATA, Ontario, Canada K2K 2E2
Telephone: (613) 592-3030
Facsimile: (613) 592-2333

AI System Development Tools

"Q&A"
Database/word processing package with an extensive Natural Language front end. Allows development of databases and documents using natural language. Can interface with popular 4th generation systems such as dBASE, WordStar, and Lotus 1-2-3. Over 50,000 copies sold in the U.S. since its introduction 2 years ago, $430.

Time Line
A powerful project management software tool that effectively manages your time, people and costs. The number of resources and dependencies allowed is limited only by memory size. Wide array of reports available, and direct interface with popular spreadsheets, and databases, $510. Time Line Graphics—add-on plotting program, $180.

NeuralWorks Professional
Sophisticated, flexible tool design to make the prototyping, training, testing and debugging of neural computing networks easy. Runs on IBM PC/XT or AT and compatibles. Provides an in-depth look at several neural computing techniques and their applications, $670.

Nexpert
Full scale expert system development tool for PCs with rich graphics support. Object hierarchy with multiple inheritance. Rules support pattern matching, integrated forward and backward chaining, automatic goal generation, non-monotonic reasoning. Open architecture allows integration with external programs. For use on IBM AT, 386, PS/2 and Macintosh II, $6425. For use on DEC VAX 2000, II, III, under VMS/UIS, $10,200.

Natural Language Query
NLQ translates English requests into the query language of a database management system, downloads the data from the mainframe to a PC, presents the data in tabular format and accommodates using the data with other PC application programs. Runs on IBM PC with 512K memory and asynchronous communication port. Currently interfaces to ORACLE and DM. Soon with DB2 and other SQL-based DBMSs. VAX and CDC Database Management Systems hosts are supported. $650.

Savvy PC
Sophisticated applications generator which has a powerful pattern recognition facility. Three major components: the Savvy Database Manager to store, manipulate, and retrieve information in plain-language style; the Savvy Language & Tutorial to allow customization of database applications; and the Savvy Retriever to translate English queries using pattern recognition. Runs on IBM PC or compatibles. $310.

ACCENT SA
A high quality text-to-speech convertor converts ordinary ASCII text into intelligible speech. Can be plugged directly into any IBM PC, PC/XT, PC AT or compatibles. Can speak in two modes: the text mode and the spell mode. Speech output may be directly connected to a speaker. A standard RS-232C serial port is provided for interface with a computer or terminal as the host. Prices range from $505—$1480.

LPA PROLOG Professional
Prolog and Expert System prototyping tool best suited for scientific and educational applications. Refined syntax gives flexibility appropriate for sophisticated programming. Several thousand copies sold throughout the world. PC Compiler version, $1210; PC Interpreter version, $535; Apes Tool Kit, $675; LPA MacPROLOG, $670.

Arity Knowledge Systems Development Pack
A new generation AI prototyping tool for PCs. A prolog inerpreter/compiler, a rule/frame-based Expert Systems shell with object hierarchy and inheritance, the external file and database access, a screen design kit, linkage to conventional languages, and an IBM SQL access, $1430.

Xi Plus
A best-selling expert system development tool from Europe. Features include: automatic generation of menus from the knowledge base itself, both forward and backward reasoning with inferencing controls from the knowledge base, help facility to get context related information at any time, and flexibility to use Xi Plus with menus and commands. $1650.

Smalltalk/V
Pure object oriented programming. Provides access to a rich tool set: bit-mapping graphics, a point-and-click user interface, multiple windows, and a source-level debugger. Runs on IBM PC/XT/AT/PS or compatibles with at least 512K memory and one of the following graphics boards: Hercules, CGA, EGA, VGA, AT&T 6300, Toshiba T3100 or IBM 3270. Add on features include: EGA/VGA Color Extension Pack, Communications Application Pack, and "Goodies" Application Pack. Price ranges from $130—$325.

Dragon Systems, Inc.
Dragon's VoiceScribe and DragonWriter voice recognition systems have been acclaimed the best voice recognizers for IBM XT, AT and 386 machines. They have constantly outperformed systems costing several times more. The design is based on a sound set of algorithms developed by two of the top researchers in the research community, Drs. Jim and Janet Baker. The devices are ideal for such applications as command/control, data entry/retrieval, form filling (hospital, government), documentation creation or dictation, front-ending of expert systems, and support systems for the handicapped. From $1,300 to $5,900.

ENQUIRY AND ORDER: Call (613) 592-3030 or send Purchase Order to Applied AI Systems, Inc.
Shipping and Handling charges and provincial sales tax (Ontario residents only) apply.

Applied AI Systems, Inc. is an authorized dealer of the above AI software/hardware products.
Prices are subject to change without notice.

Applied AI Systems, Inc.
Gateway Business Park
300 March Road, Suite 602
KANATA, Ontario, Canada K2K 2E2
Telephone: (613) 592-3030
Facsimile: (613) 592-2333