



Canadian Artificial Intelligence Intelligence Artificielle au Canada

April/avril 1988

No. 15

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

Software Liability: Are Expert Systems Any Different?

Graham Douglas Mackintosh

Droit et Logiciel: Les systèmes experts sont-ils différents?

Conference '88 Program

La programme de la Conférence '88

MAPLES and EVES at I.P. Sharp Associates (Ottawa)

Graeme Jones and/et Bill Pase

MAPLES et EVES à I.P. Sharp Associates (Ottawa)

Intelligent Tutoring Systems:

The HERON Project

Claude Frasson

Systèmes d'Enseignement Intelligemment Assisté par Ordinateur:

Le projet HERON

Artificial Intelligence Research and Applications at the Alberta Research Council

Sheila McIlraith

**Recherches et applications en intelligence artificielle
au Conseil de Recherche de l'Alberta**

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, GoldWorks and Golden Common LISP Developer are trademarks of Gold Hill Computers, Inc. Copyright 1988, Gold Hill Computers, Inc.

Gold Hill Computers, Inc.

26 Landsdowne Street
Cambridge, MA 02139

Toll-free: 1-800-242-LISP
In Mass.: (617) 621-3300

Contents

Communications

- Executive Notes 3
- Notes from Members 4
- Letters to the Editors 4

AI News

- Short Takes 7
- AI Engineering Issues Provoke Debate 12
- Connie Bryson*

Feature Articles

- Software Liability: Are Expert Systems Any Different? 13
- Graham Douglas Mackintosh*
- Conference '88 Program 15

Research Reports

- MAPLES and EVES at I.P. Sharp Associates (Ottawa) 26
- Graeme Jones and Bill Pase*
- Intelligent Tutoring Systems: 27
- the HERON Project
- Claude Frasson*
- Artificial Intelligence Research and Applications 29
- at the Alberta Research Council
- Edited by Sheila McIlraith*

Conference Reports

- Compint '87 at Montech '87: 39
- The Computer Aided Technologies
- Suhayya Abu-Hakima*
- The Second Conference on 42
- Applied Natural Language Processing
- Jugal Kalita*

Publications

- Book Reviews 45
- Books Received 48
- Computational Intelligence Abstracts* 49
- Conference Announcements 55**

Communications

- Notes administratives
- Notes pour les membres
- Lettres aux éditeurs

Nouvelles de l'IA

- Prise de vue
- Débat sur l'utilisation de l'IA en génie
- Connie Bryson*

Gros Titres

- Droit et Logiciel: Les systèmes experts sont-ils différents? 13
- Graham Douglas Mackintosh*
- Programme de la Conférence '88 15

Rapports de Recherches

- MAPLES et EVES à I.P. Sharp Associates (Ottawa) 26
- Graeme Jones et Bill Pase*
- Systèmes d'Enseignement Intelligemment 27
- Assisté par Ordinateur: le projet HERON
- Claude Frasson*
- Recherches et applications en intelligence artificielle 29
- au Conseil de Recherche de l'Alberta
- Édité par Sheila McIlraith*

Rapports des Conférences

- Compint '87 à Montech '87: 39
- Les technologies assistées par ordinateur
- Suhayya Abu-Hakima*
- La deuxième conférence au sujet du 42
- traitement pratique de langues naturelles
- Jugal Kalita*

Publications

- Critiques de livres
- Livres reçus
- Résumés d'*Intelligence informatique*

Announcements des Conférences

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'adresse à:

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

President/Président: Dick Peacocke, Bell-Northern Research, Box 3511, Station C, Ottawa, Ont K1Y 4H7.

Phone: 613-765-2629. BITNET: richard@bnr.netnorth

Vice-President/Vice-Président: Renato De Mori, School of Computer Science, McGill University, Montréal, Qué H3A 2K6.

Phone: 514-398-7072. UUCP: renato@musocs

Secretary/Secrétaire: Bill Havens, Tektronix Research Labs MS-50-662, PO Box 500, Beaverton, OR 97077, U.S.A.

Phone: 503-627-5151. CSNET: havens@crl.tek.com

Treasurer/Trésorier: Randy Goebel, Department of Computing Science, University of Alberta, Edmonton, Alta T6G 2H1.

Phone: 403-432-2683. UUCP: goebel@alberta

Canadian Artificial Intelligence

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.

Advertising:

Advertising rates and press kits are available upon request from the address below, or by phoning 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

Intelligence Artificielle au Canada

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* apparaît 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.

Réclame:

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

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

Senior Editor/Redacteur: Marlene Jones
Editor/Editeur: Sheila McIlraith
Editor Emeritus/Redacteur emeritus:
Graeme Hirst
Assistant/Aide:
Farran Sloan
Translation/Traduction:
Jean-Pierre Corriveau, Yves Lespérance,
Lem MacDonald
Section Editors/Editeurs des sections:
Advertising/Réclame:
Richard Kasperski, Margaret Mendoza
Articles/Articles: Tony Copping,
Ken Gamble, Kevin Wipond
Conferences/Conférences: Julia Driver,
Christopher Prince, Chris Stang
Humour/Humor: Rob Aiken
News/Nouvelles: Lynn Sutherland
Publications/Publications:
Graeme Hirst, Ruby Loo

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

Copyright © 1988, Canadian Society for Computational Studies of Intelligence. All rights reserved; *Canadian Artificial Intelligence* may not be reproduced in any form without the written permission of the editors. Printed in Canada by Commercial Graphics Limited. *Canadian Artificial Intelligence* is published with the assistance of the Alberta Research Council. The opinions expressed herein are those of their respective authors and are not necessarily those of their employers, CSCSI, *Canadian Artificial Intelligence*, the editors, CIPS, or the Alberta Research Council.

Copyright © 1988, Société canadienne pour l'étude de l'intelligence par ordinateur. Tout droit réservé; *Intelligence artificielle au Canada* ne doit être reproduite par quelque moyen que ce soit sans le consentement écrit des éditeurs. Imprimée au Canada par Commercial Graphics Ltée. *Intelligence artificielle au Canada* est publiée avec l'aide du Conseil de Recherche de l'Alberta. Les opinions exprimées dans ce magazine sont celles de leurs auteurs respectifs et non pas nécessairement celles de leurs employeurs, de la SCEIO, de l'*Intelligence artificielle au Canada*, des éditeurs, de l'Association canadienne informatique, ou du Conseil de Recherche de l'Alberta.

COMMUNICATIONS

Executive Notes

Once again Sheila and I are faced with the task of pulling together the various submissions and advertisements to create an issue of *Canadian Artificial Intelligence*, and I ask myself how can we convince more readers to provide us with material — small or large amounts. What do we have to do to convince you that your contribution is critical — whether it be a brief research report, an opinion piece, the abstracts of your recent publications, a conference report, or an announcement? We do need to hear from you. We aren't expecting each of you to submit material on a quarterly basis. After all, we couldn't possibly handle a thousand contributions! But if each reader simply sent us one contribution every year or two, that would suffice nicely! Some of you, of course, submit more often than that, but well, you get the picture.

So, what are you doing May 1-14? Like most of us, you will probably be trying to squeeze too much into a too short time frame and an overly tight schedule. But we would still like to convince you to squeeze in one more item — a contribution to *Canadian AI*. If the May 15 deadline is really impossible, don't despair; there will be two further possibilities this year. The deadline for the October issue is August 15 and the deadline for the January '89 issue (I can't believe that I'm planning this far ahead!) is November 15. In particular, we want conference reports for the following conferences (among others):

- Intelligent Tutoring Systems, June 1-3, Montreal
- CSCSI'88, June 6-10, Edmonton
- AAAI, August 22-26, St. Paul, Minnesota

Any volunteers? We also want opinion pieces (pick your favourite topic), research reports, etc. So please contribute!

This is also the last editorial that I plan to write this year. I'm hoping the members will play a role in writing editorials for the magazine. In particular, I'm hoping that the CSCSI executive will be particularly active in this regard. We're off to a good start for the July '88 issue; **Bill Havens** has already agreed to write an editorial on "AI and Industry".

While I'm on the topic of volunteers, it is time to consider the slate of officers for the next CSCSI executive. In particular, the position of treasurer will soon be vacated by **Randy Goebel**. Randy has done a superb job as treasurer and whoever assumes the role next has the distinct advantage of inheriting a very organized set of files. Sheila and I would also like to thank Randy for answering our many queries regarding the running of the society and financial matters related to the magazine. If you've been procrastinating regarding taking an active role in the running of CSCSI, this is the time to step forward.

There is one other matter regarding the running of the society which should be raised — memberships. When one joins CSCSI, the information is processed through CIPS which then, upon request, provides us with labels for mailing *Canadian AI*, etc. In the short period of time that Sheila and I have been editing the magazine, two people have come and gone in the CIPS' role of membership liaison. Not surprisingly, some difficulties have arisen as a result. If you have not received a membership renewal notice, or if you know of someone who has not received a copy of the magazine, be persistent. Although we cannot remedy the situation at this end, if you are having difficulties getting your membership matters straightened out, bring it to the attention of the CSCSI executive. If it is affecting a significant portion of our membership, both CSCSI and CIPS want to know.

In conclusion, Sheila and I would like to thank a few people who continue to facilitate the publication of this magazine. Special thanks to **Jean-Pierre Corriveau** (translation), **Yves Lespérance** (translation), **Meg Mendoza** (advertising) and **Farran Sloan** (production of the magazine). We would like to thank Dr. Brian Barge of the Advanced Technologies Department, Alberta Research Council for again providing financial assistance.

Marlene Jones
Senior Editor

Call for Nominations

The CSCSI/SCEIO constitution requires that a slate of candidates be proposed for the society's next executive, to serve from 1988 to 1990. Three members of the current executive have agreed to serve a second term:

President: Dick Peacocke, Bell-Northern Research
Vice-President: Renato De Mori, McGill University
Secretary: Bill Havens, Tektronix Research Labs

No nominations have been received for the position of treasurer. If there exist multiple nominations for a position, a mail ballot of the membership will be carried out. However, if only one nomination is received for each position, the candidates will be elected by acclamation. Results will be announced at the CSCSI/SCEIO general meeting at CSCSI'88 Conference and will be published in the July '88 issue of Canadian Artificial Intelligence.

Any nominee for an executive position in CSCSI/SCEIO must be a CSCSI/SCEIO member and must have agreed to let his or her name stand. People in executive positions serve for a two-year term, commencing June 1. Nominations should be sent to:

Dr. Dick Peacocke,
President, CSCSI/SCEIO,
Bell-Northern Research
Box 3511, Station C
Ottawa, Ontario K1Y 4H7

1988 Annual General Meeting

The 1988 annual general meeting of CSCSI/SCEIO will be held 4:30 p.m., Thursday, June 9, at the CSCSI'88 Conference in Edmonton. Location will be announced at the conference.

Mises en candidature

La constitution de la CSCSI/SCEIO requiert des candidatures pour un nouvel exécutif qui servira de 1988 à 1990. Trois membres actuels ont accepté un deuxième mandat:

Président: Dick Peacocke, Bell-Northern Research
Vice-Président: Renato De Mori, Université McGill
Secrétaire: Bill Havens, Tektronix Research Labs

Aucune nomination n'a été reçue pour le poste de trésorier. En cas de multiples nominations, il y a consultation postale. Par contre, une mise en candidature unique permet d'élire par acclamation. Les résultats seront annoncés à l'assemblée générale à la conférence annuelle de la société ainsi que dans le numéro de juillet.

Tout candidat à l'exécutif doit être membre et doit avoir accepté sa nomination. Les mandats sont de deux ans commençant le 1er juin. Les mises en candidatures doivent être envoyées à:

Dr. Dick Peacocke,
Président, CSCSI/SCEIO,
Bell-Northern Research
Box 3511, Station C
Ottawa, Ontario K1Y 4H7

Assemblée générale de 1988

L'assemblée générale de la CSCSI/SCEIO se tiendra à 1630h, jeudi, le 9 juin à la conférence d'Edmonton. L'endroit précis sera annoncé à la conférence.

Correction

In our January 1988 issue, we incorrectly identified the author of the report on the Second Workshop on AI in Environmental Science as Max Kranse. In fact, the author's name is Max Krause. We apologize to Mr. Krause for this error.

Correction

Dans notre édition de janvier 1988, nous avons identifié incorrectement l'auteur du rapport sur le Deuxième colloque sur l'IA et les sciences de l'environnement sous le nom de Max Kranse. Le nom correct de l'auteur est Max Krause. Nous nous excusons auprès du M. Krause pour cette erreur.

Notes from Members

New Binding

Timothy Bult of Bell-Northern Research, Ottawa, to MacDonald-Dettwiler Associates, Richmond, B.C.

Letters to the Editors

Dear Dr. Jones:

I want to compliment you and Graeme Hirst on the CSCSI's magazine, *Canadian Artificial Intelligence*. Over the years, I've seen it grow and mature into a well balanced publication covering news items, feature articles and conference reporting. The new magazine format is quite attractive. At our magazine, the magazine format has evolved from its original production on a VAX to a desktop publishing system on the MAC. We would be more than willing to share our production trials and tribulations with your staff.

Last fall, I particularly enjoyed reading a review of AAAI-87 in Seattle. The section which reviewed the Technical Paper Sessions was comprehensive and thorough. I rarely find that type of thorough review in AI publications. I hope you can continue to provide this type of reporting to your readers.

Congratulations again on a super publication!

Sincerely,

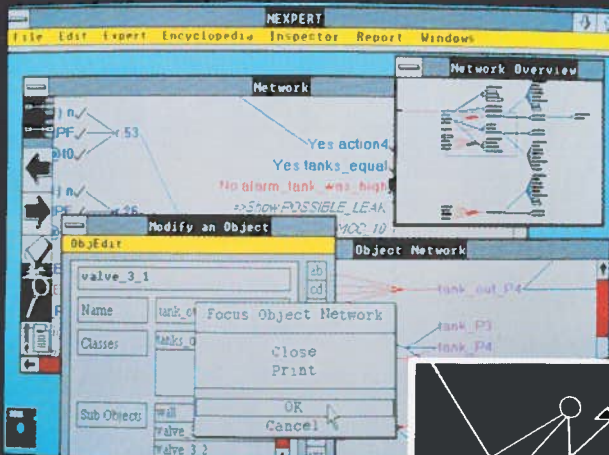
Claudia C. Mazzetti
Executive Director

American Association for Artificial Intelligence
Menlo Park, California, USA

Deadline for the
July issue is
15 May

NEXPERTTM OBJECT

The AI Library in C for Industrial Integration




The NEXPERT Runtime Library provides the solution to Embedded Delivery Systems. It is designed for complete integration of applications into operational environments written in C, Fortran, ADA, Cobol, Pascal and others.

Development & Delivery on VAX, AT, PS/2, 386, MAC II and Unix Workstations with full cross compatibility.

■ A HYBRID EXPERT SYSTEM BUILDING TOOL

■ **WRITTEN IN "C"** ■ **RULES** □ Pattern-Matching □ Integrated Forward and Backward Chaining ■ **OBJECTS** □ Methods □ Inheritance ■ **CONTROLLED INFERENCE ENGINE** □ Event Driven □ Non-Monotonic ■ **BECHTEL** □ Offers Consulting and Training Call 415/768-1500 ■ **DEC** □ NEXPERT is a Digital Distributed Software with DEC Support and Education □ Contact Your Local Digital Representative ■ **HARDWARE** □ MACINTOSH SE, II □ IBM AT, PS/2, 386 under Windows DOS, OS/2 □ \$5,000 □ DEC VAX 2000, II, III, under VMS/UIS □ \$8,000 □ Unix Workstations under X Windows □ All Cross Compatible □ For More Information Contact Neuron Data □ 444 High St. □ Palo Alto, CA 94301 □ 415/321-4488 ■ **CLIENTS** □ Leading Corporations Using NEXPERT for **DEVELOPMENT & DELIVERY**: Boeing □ NASA □ Digital Equipment Corp. □ IBM □ Shell Oil □ E. I. DuPont □ Bechtel □ Arthur Andersen □ Hughes Aircraft □ McDonnell Douglas □ Rockwell International □ G.T.E. □ AT&T □ Philips □ Polaroid □ Pect Marwick □ The Equitable □ Martin Marietta □ Eli Lilly □ EDS □ GE □ Lockheed □ TRW □ U.S. Army □ U.S. Air Force □ Telenet □ Comsat □ Northern Telecom □ Xerox □ Siemens □ Combustion Engineering □ Armco □ Teledyne □ Honeywell □ General Dynamics □ Merrill Lynch □ The Rand Corporation...



CALL-OUT

CALL-IN

PROCESSES

MANUFACTURING

CONFIGURATION AND DESIGN

TRADING

UNDERWRITING

DIAGNOSIS AND TESTING

OFFICE AUTOMATION

PROCESS CONTROL


RELATIONAL DATABASES: RDB, SQL, DBIII

ND
NEURON DATA

NeuralWorks Professional

NeuralWorks Professional (TM) is a tool designed to make the prototyping, training, testing and debugging of neural networks easy. The PC-based tool is geared to the professional engineer, as well as to scientists interested in the field. It is designed to do concept development and testing, and includes an overview of neural computing as well as an in depth look at several major paradigms. It sports an easy to use graphically oriented interface, with pop-up menus.

Applications: Some of the areas where concept testing can be done include: Signal Processing, Digital Filters, Matched Filters, Image Recognition, NP-complete problems, Fuzzy Expert Systems (Fuzzy Cognitive Maps), and many other applications of pattern recognition.



Network Specification

Paradigms Supported: Perceptron, Adaline, Madaline, Hopfield, Brain-State-in-a-Box, Bi-directional Associative Memory, Back-Propagation, Counter-Propagation, Hopfield/TSP. Other paradigms are under development.


Summation Functions: Sum, Incremental Sum, Min, Max, Majority, Product, Norm1, NormN.

Transfer Functions: Sign, Psi (BAM), Step Function, Perceptron, Linear, Bounded-Linear, Sigmoid, Hyperbolic Tangent, Sine.

Output Functions: Direct Competitive with one or two winners, Competitive Adaline Adaption.

Learning Rules: Hebb, Hebb — Anti-Hebb, Hopfield, BSB-Hopfield, BSB-Widrow-Hoff, Widrow-Hoff, Kohonen, Normalized Kohonen, Generalized Delta (Back-Propagation), Adaline, Madaline, Perceptron. Other learning rules are under development.

Special Features: Network or Layer-specific temperature schedule for simulated-annealing; learning schedules for variable learning rates across the network and other special effects; disable individual processing elements; fix weights; user-defined i/o function; and more.



System Requirements:

IBM PC, XT or AT or 100% compatible with 512K of memory (640K recommended); CGA or EGA Color, or Hercules Monochrome Graphics Adapter; MS-DOS or PC-DOS version 2.0 through 3.x; Hard disk, Numeric co-processor, and a Mouse recommended.

Price: \$670.00

Enquiry and Order: Call (613) 592-0084 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 NeuralWare products.



Applied AI Systems Inc.
4015 Carling Avenue,
Kanata, Ontario, Canada K2K 2A3
Telephone: (613) 592-0084
Facsimile: (613) 592-2333

NeuralWorks Professional is a registered trademark for NeuralWare Inc.

Short Takes

PRECARN — Precompetitive Applied Research Network

The following announcement appeared on AIList earlier this year. Deadline for submission of RFPs is April 18, 1988.

PRECARN is a not-for-profit consortium of over 30 Canadian corporations. PRECARN (sometimes expanded as Precompetitive Applied Research Network) has as its overriding objective the increase of industrial competitiveness in Canada. This will be accomplished by funding applied research projects to raise awareness, and exploitation, of artificial intelligence and robotics technologies within industry.

PRECARN is now commencing the process that will identify, and eventually fund, the research and development activities in artificial intelligence and robotics that hold the greatest promise for eventual exploitation by the member corporations. PRECARN is inviting applications for the funding of feasibility studies. Applications to this RFP may come from:

- a) Canadian universities or colleges
- b) Canadian corporations or subsidiaries of foreign corporations having, or actively moving towards, a significant research and development operation in Canada
- c) Crown corporations with an arm's-length relationship with the government.

For more information, a complete copy of the RFP, and application forms, contact: PRECARN Associates Inc., 30 Colonnade Rd., Suite 300, Nepean, Ontario K2E 7J6. Phone: (613) 727-9576

NRC opens Industrial Technology Institute

In November the Honourable Frank Oberle, Minister of State for Science and Technology, announced the opening of the National Research Council's Canadian Institute of Industrial Technology (CIIT) in Winnipeg.

The new institute will use the resources of the NRC to assist companies in developing new products and techniques in the areas of expert systems, robotics, and computer-integrated manufacturing.

"The Canadian Institute of Industrial Technology will help strengthen Canadian industry's competitiveness in world markets by assisting them in applying the latest in advanced technologies to increase their productivity," said Dr. Larkin Kerwin, President of NRC.

CIIT will provide space and services at cost to companies involved in private sector initiated projects. The \$26 million complex in downtown Winnipeg houses more than 17,000 square metres of laboratory, workshop, and office space. NRC oversees the management of the

Prise de vue

PRECARN — Le réseau préconcurrentiel de recherche appliquée

L'avis suivant a paru sur "l'AIList" plus tôt cette année. L'échéance pour la soumission de demandes des propositions est le 18 avril, 1988.

PRECARN est un consortium sans but lucratif de plus de 30 corporations canadiennes. PRECARN (un acronyme de Pre-Competitive Applied Research Network) a pour objectif primordial l'augmentation de la compétitivité industrielle canadienne. Ils comptent atteindre cet objectif en finançant des projets de recherche appliquée qui feront connaître les technologies de l'intelligence artificielle et de la robotique et susciteront leur exploitation par l'industrie.

PRECARN est maintenant en train d'initier le processus qui identifiera et financera éventuellement les activités de recherche et développement en intelligence artificielle et en robotique les plus prometteuses en vue d'une utilisation éventuelle par les compagnies membres. PRECARN invite maintenant des demandes des propositions pour des études de faisabilités. Les demandes pour ce programme des propositions peuvent venir:

- a) d'universités ou collèges canadiens
- b) de corporations canadiennes ou de filiales de corporations étrangères poursuivant ou ayant l'intention de poursuivre des activités de recherche et développement significatives au Canada
- c) des compagnies de la Couronne ayant une relation "arm's length" du gouvernement.

Pour plus d'informations, une copie complète des règlements du programme, et des formulaires de demande, contacter: PRECARN Associates Inc., 30 Colonnade Rd., Suite 300, Nepean, Ontario K2E 7J6. Phone: (613) 727-9576

Le CNR crée l'Institut canadien de technologie industrielle

En novembre l'Honorable Frank Oberle, Ministre d'état (science et technologie) a annoncé l'ouverture de l'Institut canadien de technologie industrielle (ICTI) du Conseil national de recherches à Winnipeg.

Le nouvel institut utilisera les ressources du CNR pour assister les compagnies qui développent de nouveaux produits dans les domaines des systèmes experts, de la robotique et de la fabrication intégrée par ordinateur.

"L'Institut canadien de technologie industrielle aidera à renforcer la compétitivité de l'industrie canadienne dans les marchés mondiaux en assistant celles-ci dans l'application des plus récentes technologies avancées pour augmenter leur productivité", dit le Dr. Larkin Kerwin, président du CNR.

L'ICTI offrira des locaux et des services au prix coutant aux compagnies privées qui poursuivent des projets. Le complexe de 26,000,000\$ au centre-ville de Winnipeg couvre plus de 17,000 mètres carrés de laboratoires,

institute with an operating budget of \$1.2 million a year and a core staff of 25. The CIIT has 21 industrial and educational partners occupying 60 percent of the facility.

CATA concerned about Free Trade, Competitiveness, and Tax Reform

Canada is falling behind other countries in the race for technology leadership. In a statement issued by the Canadian Advanced Technology Association, a concern over mixed signals in tax reform was expressed.

The Canadian Advanced Technology Association is a national industry association of companies having production or R&D capability located in Canada, as well as research, financial, and academic institutions, companies, or individuals, concerned with the performance of the private sector. The Association is dispirited by proposals in the White Paper on Tax Reform which limit the use of R&D investment tax credits to 50% of tax otherwise payable. This, said Mr. Roy Woodbridge, president of the association, will leave "over 40% of the research performed by the most research-driven companies in Canada, unsupported by the tax system."

The agreement in principle on free trade between Canada and the United States means that "it is now time to raise the task of building national competitiveness strength to the top of the public policy agenda." Mr. Woodbridge outlined elements of a national strategy to make Canada one of the most technologically sophisticated and competitive countries in the world. These initiatives would include:

- changes in the way in which we educate and train Canadians
- efforts to build high tech industry
- the more rapid diffusion of productivity-enhancing technology throughout the economy
- increasing public awareness of the strategic contribution that science and technology make to the well-being of Canada

Neutron Inc. Software Key Component in Australia's Automated Banking

Toronto's Neutron Inc. has been revealed as a key component in Australia's "Core System for the Nineties." The \$114 million project is to give Westpac Banking Corporation, Australia's largest banking group, the most sophisticated automated banking systems in the world. Frame-based technology will be used to capture banking knowledge in an adaptable/reusable form.

Computing Canada goes Online

The *Computing Canada* newspaper has started an on-line service for its subscribers. The service provides a method of delivering news, public-domain and share-ware software, and demonstration programs. It also allows callers to discuss management problems and solutions. *Computing Canada Online* started operation in

ateliers et espaces à bureaux. Le CNR contrôle la gestion de l'institut avec un budget d'opération de 1,200,000\$ par an et un personnel de base de 25. L'ICTI a 21 partenaires, des corporations et des maisons d'enseignement, qui occupent 60% des locaux.

L'ACTA s'inquiète du libre échange, de notre compétitivité et de la réforme fiscale

Le Canada perd du terrain face aux autres pays dans la course au leadership technologique. Dans un communiqué récent, l'Association canadienne de technologie avancée exprime sa nervosité face aux signaux contradictoires contenus dans la réforme fiscale.

L'Association canadienne de technologie avancée est une association nationale de compagnies ayant des capacités de productions ou de recherche et développement au Canada, ainsi que des institutions de recherche, des institutions financières, des maisons d'enseignement et des individus concernés par la performance de le secteur privé. L'association est déçue par les propositions du livre blanc sur la réforme fiscale qui limitent l'utilisation des crédits d'impôt pour les investissements en recherche et développement à 50% des impôts payables autrement. Ceci, dit M. Roy Woodbridge, président de l'association, laissera "plus de 40% de la recherche faite par les compagnies les plus axées sur la recherche au Canada sans support fiscal."

L'accord de principe sur le libre échange entre le Canada et les Etats-Unis veut dire "qu'on doit maintenant porter la tâche d'augmenter la force compétitive nationale au sommet de l'agenda politique." M. Woodbridge a ébauché les éléments d'une stratégie visant à faire du Canada un des pays les plus compétitifs et technologiquement avancés du monde. Ces initiatives incluraient:

- des changements dans la façon d'éduquer et de former les canadiens
- des efforts pour développer l'industrie de haute technologie
- une diffusion plus rapide des technologies pouvant augmenter la productivité dans toute l'économie
- des efforts pour favoriser la prise de conscience par le public de la contribution stratégique de la science et la technologie au bien-être du Canada

Logiciel de Netron Inc. sera une composante clé du système bancaire automatisé d'Australie

On a annoncé que Netron Inc. de Toronto produira une composante clé du "Système de base pour les années 90" d'Australie. Le projet de 114 million de dollars donnera à la Westpac Banking Corp., le groupe bancaire le plus important d'Australie, le système bancaire automatisé le plus sophistiqué au monde. La technologie des cadres sera utilisée pour capturer les connaissances bancaires dans une forme adaptable et réutilisable.

Computing Canada se branche

Le journal *Computing Canada* offre l'accès direct à ses abonnés. Le service permet de transmettre nouvelles, logiciels et démonstrations. Ceux qui appellent pourront aussi discuter gestion du journal, qui offre ce service

November.

McGill, IBM in 3 year, \$5 million Project

McGill University and IBM Canada Ltd. have entered into a project to develop advanced computer applications, expert systems and educational software. IBM will provide more than 150 PS/2 computers and software. McGill will contribute faculty research time and facilities. One area of investigation will be adding computer applications to the research performed by McGill's faculty of Medicine and by its Cognitive Sciences Centre in the areas of language comprehension and learning disorders.

Canadian Commercial Expert System

D&S Knowledge Systems Inc., a Calgary firm, has announced the successful completion of a practical knowledge-based software system which has been used as an "assistant" in the interpretation of well logs. The product is the result of a two year, \$2 million joint venture with the Alberta Research Council.

The product integrates a knowledge base with log analysis software. The system is "open", so that a user may change the knowledge base, or add in knowledge of new procedures. The software runs on personal computers and may be modified for other domains. D&S is offering its services in consulting on Expert Systems applications and log analysis. For further information, contact: Gerald Parks, Manager, Business Development, D&S Knowledge Systems Inc. Phone: (403) 268-6500.

NRC Publishes Research Directory

The National Research Council has released a Directory of Research Activities, a compendium of information on its current scientific programs, capabilities, facilities and services. The directory (NRC No. 27961) is available for \$10 in English or French from Publications Sales and Distribution, Bldg. R-88, National Research Council Ottawa, Canada K1A 0R6.

Toronto Company Offering Super-Computing Services

Revelations Resources Ltd. of Toronto are offering supercomputing services at affordable prices. A Convex Computer Corporation C1-XP will be accessible through dial-up facilities or in a client workroom in Mississauga. An extensive technical library, and educational and technical support are available. For further information, contact: Ian MacLachlan, Marketing Director. Phone: (416) 270-6667.

Myrias in Texas

Myrias Research Corp., an Edmonton based developer of parallel computers has set up a subsidiary company in Houston to serve the U.S. market. The company has received substantial private and government interest in its new parallel supercomputer.

depuis novembre.

McGill et IBM ensemble pour un projet de 5 millions dollars sur trois ans

L'université McGill et IBM Canada Ltd. se sont lancés dans un projet visant à développer des applications de pointe, des systèmes experts et des logiciels éducatifs. IBM fournira plus de 150 PS/2 ainsi que le logiciel et McGill, le temps de travail et les autres ressources. On compte entre autres créer des applications en compréhension linguistique et en difficultés d'apprentissage pour les professeurs de la faculté de médecine et du centre des études cognitives de l'université.

Système expert commercial canadien

D&S Knowledge Systems Inc. de Calgary vient d'annoncer qu'il a complété avec succès un logiciel pratique à base de connaissance qui est employé à titre d'assistant pour l'interprétation des relevés de puits.

Le système est 'ouvert' en ce qu'il permet à l'utilisateur de modifier la base de connaissance ainsi que d'ajouter de nouveaux mécanismes d'interprétation. Pour plus d'informations: Gerald Parks, Manager, Business Development, D&S Knowledge Systems. Phone: (403) 268-6500.

Le CNR publie un annuaire de chercheurs

Le Conseil National de la Recherche a publié un répertoire des activités en recherche, soit une liste des programmes, potentiels, matériels et services existant. L'annuaire (NRC no. 27961) se vend à dix dollars, en français ou en anglais. Il est disponible chez: Publications Sales and Distribution, Bldg. R-88, Conseil National de la Recherche, Ottawa, Canada K1A 0R6.

Une compagnie de Toronto offre l'accès à un super-ordinateur

Revelations Resources Ltd. de Toronto offre l'accès à prix abordable à un super-ordinateur. Un C1-XP de Convex Computer Corporation pourra être rejoint à partir de lignes téléphoniques ou d'une salle de travail à Mississauga. Le service comprend l'aide technique ainsi qu'une vaste bibliothèque technique et éducationnelle. Pour plus de renseignements: Ian MacLachlan, Marketing Director. Phone: (416) 270-6667.

Myrias au Texas

Myrias Research Corp., un créateur d'ordinateurs parallèles basé à Edmonton a ouvert une filiale à Houston pour servir le marché américain. La compagnie a suscité un intérêt considérable autant du côté gouvernemental que de la part du secteur privé, en rapport avec son nouveau superordinateur parallèle.

Gigamos Sues Gensym

Gigamos Systems, Inc., of Montreal, who earlier this year bought the assets of Lisp Machines Inc., are suing Gensym Corporation. Gensym Corporation was founded by employees of LMI who left the company before it went bankrupt. The suit claims that the employees used LMI trade secrets to develop G2, an expert system tool that competes with Picon, now a Gigamos Systems product. The suit also claims that the employees left LMI at a critical moment, and actually caused LMI's bankruptcy.

OOPS Report

A new quarterly publication on object-oriented programming begins in the spring. The OOPS Report will cover reviews of OOPS developments, ranging from the human-computer interface to the new generation of object-oriented databases. Subscriptions (4 issues): UK £95, Overseas £105. Contact: The OOPS Report, PO Box 84, Plymouth PL1 1SZ, England. Phone: 0752 673814.

IBM Investigates Supercomputers

IBM has joined forces with a three-month old company, Supercomputer Systems Inc. of Eau Claire, Wisconsin to develop a new generation of supercomputers. Steven Chen, founder of SSI and a former chief engineer at Cray Research Inc., will have responsibility for the design of the supercomputing system. IBM will provide Chen with initial financing, personnel and access to its technology. Resulting products are not expected until the early 1990s.

New Products

MacII® with Explorer™ Lisp Chip

Texas Instruments and Apple Computer, Inc. have announced a value-added re-seller agreement for a new AI workstation. The microExplorer computer system combines TI's advanced Explorer™ AI microprocessor and software environment with the Apple® Macintosh® II personal computer. The microExplorer processor supports up to 12 megabytes of memory and plugs into a single NuBus™ slot in the Macintosh II. Explorer and Macintosh II operating environments execute concurrently, and are closely coupled with multi-processor software developed by TI.

List prices for the microExplorer systems start at US\$14,995. The base configuration includes a Macintosh II with two megabytes of memory, a 40-megabyte disk, microExplorer processor with four megabytes of memory, and Explorer runtime software environment. Full Explorer development software, additional memory and mass storage options are available. System deliveries are scheduled to begin in the second quarter of 1988.

Many third party software vendors have indicated their intentions to port their software to the new computer. Intellicorp's KEE®, Carnegie Groups' Knowledge Craft®, and Inference's ART®, all high-end

Gigamos poursuit Gensym

Gigamos Systems, Inc., de Montréal, qui a acheté les actifs de Lisp Machines Inc. plus tôt cette année, a intenté une poursuite à Gensym Corporation. Gensym Corp. fut fondée par des employés de LMI qui ont quitté cette compagnie lorsqu'elle a fait faillite. La poursuite prétend que les employés ont exploité des secrets d'affaires de LMI pour développer G2, un outil pour systèmes experts concurrent du produit Picon de Gigamos Systems. La poursuite prétend aussi que les employés ont quitté LMI à un moment critique et ont effectivement causé la faillite de LMI.

OOPS Report

Ce nouveau bulletin sur la programmation à base d'objets commencera à paraître au printemps. Le "OOPS Report" examinera les nouveaux développements dans le domaine, à partir de l'interface homme-machine jusqu'à la nouvelle génération de banques de données à base d'objets. Souscriptions (4 numéros): UK £95, outre-mer £105. Contacter: The OOPS Report, PO Box 84, Plymouth PL1 1SZ, England. Tél: 0752 673814.

IBM s'intéresse aux superordinateurs

IBM s'est associé à une compagnie vieille de trois mois, Supercomputer Systems Inc. d'Eau Claire, Wis., pour développer une nouvelle génération de superordinateurs. Steven Chen, le fondateur de SSI et ancien ingénieur en chef à Cray Research Inc., aura la responsabilité de la conception du système. IBM fournira à Chen du financement initial, du personnel et l'accès à sa technologie. Les produits qui en résulteront ne sont pas attendus avant le début des années 1990.

Nouveaux produits

MacII® avec puce Lisp Explorer™

Texas Instrument et Apple Computer, Inc. ont annoncé une entente de revente avec valeur-ajoutée pour un nouveau poste de travail d'IA. Le système microExplorer joint le microprocesseur avancé et l'environnement logiciel d'IA Explorer™ de TI à l'ordinateur personnel Macintosh® II de Apple®. Le processeur microExplorer supporte jusqu'à 12 mégaoctets de mémoire et se branche sur une seule fiche NuBus™ du MacII. Les environnements d'exploitation de l'Explorer et du MacII exécutent concurremment et sont étroitement associés par un logiciel de multi-processeur développé par TI.

Les prix de liste pour le système microExplorer commencent à 14,995\$US. La configuration de base comprend un MacII avec 2 mégaoctets de mémoire, un disque de 40 mégaoctets, le processeur microExplorer avec 4 mégaoctets de mémoire, et l'environnement logiciel temps d'exécution Explorer. Le logiciel de développement complet de l'Explorer et des options de mémoire et de mémoire de masse additionnelle sont aussi disponibles. Les premières livraisons sont prévues pour le deuxième trimestre de 1988.

Plusieurs vendeurs de logiciel indépendants ont indiqué leur intention de transporter leur logiciel sur le

knowledge tool-kits, have been announced. Wisdom Systems' Concept Modeller™, Gensym's G2 Real-Time Expert System, Graphael's G-Base™ object-oriented database, and ICAD Inc.'s ICAD System, have also been announced. These products are expected in mid-summer.

Butterfly Lisp

BBN Advanced Computers, Inc. has announced the availability of a shared memory multiprocessor Common Lisp called Butterfly Lisp. Butterfly Lisp runs on the Butterfly multiprocessor. It supports both Common Lisp and Scheme, both interpreted and compiled.

Logicware's TWAICE on 386

The TWAICE expert system shell, available since 1985 on IBM mainframes, DEC VAX and M68000 workstations, is now available on 386-based hardware. Logicware's commitment is to provide expert system development tools to corporate MIS and DP departments. Its other products include MProlog and the AI EDGE, a prototyping service.

KEE on 386

Intellicorp's Knowledge Engineering Environment™ is now available on 386 PCs. The recommended configuration is 10 megabytes of main memory and a 100 megabyte hard disk. The system includes Common Lisp and Microsoft Windows™ and is available under DOS or UNIX™. The software is available at an introductory price of US\$9900.

Hypercard and CD ROM for Macintosh

Apple is nearing release of a CD ROM drive for the Macintosh accompanied by a new version of Hypercard. Apple's drive is a half-height SCSI device made by Sony. It has an average access time of 500 milliseconds and will cost US\$15,500.

Prolog with Embedded C Compiler

Arity Corporation has announced the release of Arity/Prolog Version 5 which includes an embedded C compiler, allowing the use of C and Pascal in conjunction with Prolog. The compiler allows C declarations, preprocessor directives, and C expressions embedded within Prolog source code. The software runs on PCs and also includes screen design predicates to create windows, menus, and dialog boxes. The Arity/Prolog Compiler and Interpreter sells for US\$650. The interpreter alone sells for US\$295. For more information, contact: Meredith Bartlett, Arity Corporation, 30 Domino Drive, Concord, MA 01742. Phone: (617) 371-1243.

nouvel ordinateur. Le KEE® de Intellicorp, Knowledge Craft® de Carnegie Group, et le ART® de Inference, tous des ensembles d'outils de connaissance de haut niveau, ont été annoncés. Le Concept Modeller™ de Wisdom System, le G2 Real-Time Expert System de Gensym, la banque de donnée à base d'objets G-base™ de Graphael, et le ICAD System de ICAD Inc. ont aussi été annoncés. Ces produits sont attendus vers le milieu de l'été.

Butterfly Lisp

Un multiprocesseur CommonLisp à mémoire partagée appelé Butterfly Lisp est disponible chez BBN Advanced Computers, Inc. Butterfly Lisp utilise un multiprocesseur Butterfly et comprend Common Lisp et Scheme, interprétés et compilés.

Le TWAICE de Logicware sur 386

L'environnement pour systèmes experts TWAICE qui existe depuis 1985 sur les gros ordinateurs IBM, le VAX de DEC et les postes de travail utilisant le M68000, est à présent disponible sur le processeur 386. Logicware produit des logiciels de développement de systèmes experts pour les départements de gestion et de traitement de données de corporations, ainsi que MProlog et AI EDGE, un système de prototypes.

KEE sur un 386

Le Knowledge Engineering Environment d'Intellicorp est maintenant disponible sur les PC 386. On recommande 10 meg de mémoire centrale et 100 meg sur disque rigide. Le système comprend Common Lisp et Microsoft Windows et roule sous DOS et UNIX. Le logiciel est offert au prix spécial de 9900\$US.

Hypercard et CD ROM pour le MacIntosh

Apple prépare la mise en marché d'un lecteur CD ROM pour le MacIntosh. Le lecteur vient avec une nouvelle version de Hypercard et consiste en un appareil SCSI de demi-hauteur fabriqué par Sony. Le temps d'accès est de 500 millisecondes et le prix de 15,500\$US.

Prolog incluant un compilateur C

Arity Corporation a annoncé la mise en marché de leur Prolog version 5 qui comporte un compilateur C permettant de combiner C et Pascal avec Prolog. On peut utiliser les déclarations, les commandes de pré-traitement et les expressions C à l'intérieur de Prolog. Le logiciel, qui en coûte 650\$US pour le compilateur et l'interpréteur, tourne sur les PC et comprend des prédicats pour la création de fenêtres, menus et boîtes de dialogue sur l'écran. L'interpréteur seul coûte 295\$US. Pour plus de renseignements: Meredith Bartlett, Arity Corporation, 30 Domino Drive, Concord, MA 01742. Phone: (617) 371-1243.

AI Engineering Issues Provoke Debate

by Connie Bryson

Débat sur l'utilisation de l'IA en génie

RÉSUMÉ: La communication invitée de David Parnas "Pourquoi les ingénieurs ne devraient pas utiliser l'intelligence artificielle" à la conférence annuelle 1987 de l'ACI à Edmonton a provoqué des discussions passionnées. La critique de Parnas fut centrée sur l'absence de définitions claires de l'IA et de plusieurs des termes utilisés par les chercheurs en IA. Des commentaires de Randy Goebel, Renee Elio, John Mylopolous, Mark Fox et Brian Gaines ajoutèrent au débat.

For supporters of the U.S. Strategic Defence Initiative, Soviet missiles have frequently been the least of their worries. They've had more trouble from outspoken critics like David Parnas, a computer science professor from Queen's University.

However in November, at the 1987 CIPS Edmonton annual conference co-sponsored by the University of Alberta, Parnas mentioned "Star Wars" only in passing. His paper, "Why engineers should not use artificial intelligence", was aimed at a new target — artificial intelligence.

Parnas didn't mince words when it came to criticizing AI's goals and techniques. Not surprisingly, the talk generated heated discussion. Lucid reasoning was, at times, punctuated by verbal mudslinging.

At the heart of Parnas' argument was his criticism of the lack of a clear definition for AI and many of the terms used by AI researchers.

He took issue with the idea that AI is concerned with a "theory of knowledge", claiming this definition puts AI in the realm of psychology or philosophy. He said he finds nothing wrong with this area of research as long as the researchers who are involved in it stay out of computer science.

"What I do object to is when these people start telling those who should be doing engineering — producing practical systems on a day-to-day basis — that there is a technology just waiting there for them to use which is the future of the world," he said.

Randy Goebel, University of Alberta, said outrageous claims have been made for many new disciplines, and that AI was by no means the worst offender. He added that most of these claims are made by people who have no understanding of AI.

"It's important to make a distinction between people who are studying the scientific discipline of artificial intelligence and those who are just cashing in on the bandwagon," Goebel said.

Renee Elio, University of Alberta, responded to Parnas' critique of expert systems. She agreed with his statement that studying how a human expert solves a problem, then transferring that knowledge to a computer program, is the wrong way for a computer to solve problems. However, Elio pointed out that "the really successful expert systems, economically or theoretically, do not

mimic the reasoning of humans but rely on an analysis of the problem."

Both John Mylopolous, University of Toronto, and Mark Fox, Carnegie-Mellon University, were frustrated by the style of Parnas' presentation. His talk was wide-ranging, touching on many topics, such as non-monotonic logic and heuristic programming, only superficially.

"In order to have an intelligent discussion about artificial intelligence and what it tries to do, I think that there has to be a much deeper presentation of the issues," Mylopolous said.

Fox went further in his criticism.

"I find the mode in which you present this information highly destructive with regards to computer science because it attacks another field as opposed to understanding what's going on in it...I expect more from you."

Parnas denied that he was on the attack.

"I'm trying to strengthen computer science so that people learn to use proper terminology and talk about things in as precise a way as they are capable of. It's not my intention to attack a field — I don't consider AI a field," Parnas said.

Mylopolous challenged Parnas' contention that AI has a monopoly on fuzzy meanings and pointed to Parnas' own field, software engineering, for examples.

"It seems to me that artificial intelligence is no worse than some other areas like software engineering," Mylopolous said.

"Could there be something like a gray area around the term software? What about the term computer? Is a hand calculator a computer?"

The more emotional reactions to Parnas' talk drew a reflective response from Brian Gaines, University of Calgary. He found the presentation valuable because it raised questions about social responsibility.

"(Within computing) we can fly kites, have grand aspirations, and use certain vocabulary, but we know what we're talking about...It's an established part of the culture," Gaines said.

"But when that culture gets outside, it ends up sounding like we're going to solve the world's problems with some pie in the sky.

"It's very legitimate to raise these issues, we aren't just a scientific community, we're part of society. ... I think it's important to have the guts to face these issues. Some of this discussion will be very useful to the profession."

Connie Bryson is a free-lance technical writer based in Edmonton.

Software Liability: Are Expert Systems Any Different?

by Graham Douglas Mackintosh

Droit et Logiciel: Les systèmes experts sont-ils différents?

RÉSUMÉ: Les systèmes experts présentent des caractéristiques qui les distinguent des logiciels traditionnels et brouillent leur statut légal en ce qui a trait à la responsabilité d'une erreur de logiciel entraînant un préjudice. Le problème de la négligence se complique pour ces systèmes où le client affecte le comportement du logiciel. La négligence pourrait être celle de l'expert dont le système modélise plus ou moins fidèlement l'expertise. Plus les systèmes d'IA approchent de la définition sociale d'une personne, plus il devient difficile de les traiter en tant que simples outils créés par l'homme.

December, 1975: Florida state trooper Robert Rennie slowly approached a car parked on the shoulder of a highway. He had been informed by the central computer of the Florida Crime Information Centre that the vehicle was registered as stolen, and that he should exercise caution. Rennie became alarmed when the driver, Frank Booth, appeared to move aggressively; he drew his handgun and killed Booth in self-defense. It later became apparent that a software error had caused the computer to generate an incorrect report; the car had not been stolen, and Frank Booth had been an innocent man [1].

March, 1986: Voyne Ray Cox prepared himself for another session of radiation therapy in the East Texas Cancer Centre. As he lay on the treatment table, he was suddenly jolted by pain as the Therac-25 therapy system released a massive overdose of radiation into his shoulder. Six months later, after lapsing into a coma, Cox died from radiation poisoning. The malfunction has since been traced to a bug in the Therac-25 control software. Cox's widowed spouse has filed suit against the manufacturer, Atomic Energy of Canada Ltd [2].

These were real people whose lives were lost through software malfunctions. How does the legal system currently deal with liability for software failure? Should expert systems be similarly treated, or do they possess characteristics that require special consideration?

Traditional Software Liability

The word "traditional" scarcely applies, since software liability is still in a state of turmoil. Much of the debate centers on whether software is a product or a service. The issue is an important one, since the legal system treats product and service liability differently.

In the case of a product, such as a car, the manufacturer is held strictly liable for any physical damage the defective product may cause, regardless of the competence exhibited during design and construction. If a car's brakes fail due to a design flaw, the builder is liable for harm to persons and property; proof of manufacturer negligence is not required. In contrast, service liability is applicable only if there is evidence of breach of contract, or negligence in the provision of the service. A consultant may give advice that later proves to be incorrect; but liability claims would be groundless if the error lay within the bounds of professional competence.

The transient nature of software makes it difficult to classify; in some ways software seems to be a product: it resides on a tangible medium, it can be mass produced, and it can be bought and owned. In other respects it appears to be a service: it can be uniquely tailored to a client, and can offer the same interactive benefits as its human counterparts [3].

The courts have therefore avoided making a sweeping classification, and have instead proceeded on a case-by-case basis. Although this practice makes it difficult for software engineers to predict the extent of their liability, it allows the courts to operate using case law, rather than passing new legislation. The real question now becomes: Can AI systems be viewed through the same product-service framework, or do they possess attributes that make such a position untenable? First, consider the issues that complicate expert system liability beyond that of traditional software.

Self-Modification

Expert systems are becoming learning systems. In the future, they will improve and adapt their own knowledge-bases so that they can perform more expertly in their particular environments. As an extreme example of this, consider the artificial neural systems now being developed, in which self-modification is fundamental to their operation. A "veteran" neural network may differ substantially from its original state.

Since learning systems can tailor themselves to a particular client, they quickly lose the mass-produced quality that is common to products, and acquire the "one-of-a-kind" attribute that implies a service. They may be more easily exempted from strict liability, thereby forcing a plaintiff to consider manufacturer negligence.

However, a self-modifying expert system also complicates the issue of design negligence, since the manufacturer might argue that the system was well behaved when it left the factory, so the client must have "brought it up" badly. In their eyes, it is not unlike a car owner modifying the vehicle's brakes and then suing the manufacturer when they fail. How can the manufacturer be held responsible, when the client has been so involved in the system's construction?

Finally, learning expert systems may develop an understanding of their subjects that surpasses our own.

This has further implications that are discussed below.

Cloned Liability

The printing press allows information to be disseminated independently of the human mind. Accordingly, there is an explosive potential for liability, since an erroneous publication can have consequences beyond the scope of the original source. Similarly, artificial intelligence permits expertise to be used independently of the mind in which it was originally developed. Once again, there is new potential for the extension of liability. For example, a doctor is hired by a software company to provide the knowledge for an expert system on abdominal pain. Unknown to the knowledge engineers, the doctor provides incompetent advice, and by year's end, five thousand negligent "clones" of the doctor have been sold.

From the doctor's perspective, he had only one client: the software developers. Does his professional liability end there? It is instructive to notice that in some areas of the US, an accountant who negligently audits a client's books is liable for the subsequent financial loss of outside investors. In both cases, the issue is whether professional liability should be extended to persons other than the immediate client [4]. Clearly, the potential exists for "cloned" professionals to have their vulnerability to malpractice greatly magnified.

Professional Systems

Does the malfunction of a medical expert system really represent a new legal challenge? After all, traditional software systems are used extensively by physicians and can fail with equally fatal results (a fact that was sadly illustrated by the malfunction of the Therac-25 radiation therapy system). Unlike traditional software, however, an expert system attempts to emulate the doctor in his or her licensed capacity. As discussed above, this emulation can be so complete that the human on which the system is based may be liable for the system's errors. At what level of intelligence, if any, does the expert system cease to be a tool, and become a colleague? This is an important distinction, since a doctor that refers a patient to a specialist is not liable for that specialist's treatment. Similarly, a doctor that refers a patient to a "reputable" expert system should not be liable for that system's errors [5].

How does one define "reputable"? The most obvious solution is to test expert systems as we do now for human professionals, and license those that pass. Doctors can then refer their patients to expert systems without fear of liability, since the minimum professional requirements have been assured.

This leads us to wonder who would be liable in the event that computer malpractice did occur. The computer has no assets to give up in compensation, and (beyond turning the power off) cannot be punished. The injured party might turn to the manufacturer, or the licensing officials for retribution. Again the problem of self-modification is raised, with both of the defendants claiming that the system was operating correctly at the time of their inspections. Furthermore, attempts to simplify this predicament by disabling the learning mechanisms may also result in malpractice, since failure to remain aware of medical advances constitutes negligence.

A more comfortable solution may be to disallow machine licensing altogether. Expert systems, no matter how competent, are only advisors and must be treated as such. The ultimate responsibility for treatment must lie with the human professional. We could then rely on the knowledge of the human doctor, while receiving the luxury of a second opinion from the expert system.

This solution becomes less tenable when we remember that self-modifying expert systems may eventually outperform their human counterparts. The value of human professionals would become dubious, since they would not dare over-rule the computer's judgment for fear of being charged with negligence. Furthermore, the justice of holding the human professional responsible for treatment, in spite of the computer's abilities, becomes increasingly questionable. Is it not reasonable for them to develop a reliance on the expert system's judgment? Consider the ability of today's computers to accurately, and consistently, tally a column of numbers. Our faith in them has become so great, that manual verification of their work would be considered a waste of effort. If an obscure bug finally caused a computer to incorrectly add a list of numbers, would it be reasonable to blame the operator for not checking the addition?

Conclusion

Even in their infancy, expert systems possess characteristics that set them apart from traditional software and confuse their legal standing. There can be little doubt that future developments in the field of artificial intelligence will further complicate an already turbulent legal situation. Perhaps the legal system's difficulties stem from a single source: the closer AI systems come to the social definition of a person, the harder it becomes to treat them as creations of human industry.

References

- [1] James E. O'Connor, "Computer Professionals: The Need for State Licensing", *Jurimetrics Journal*, Spring 1978, p. 256-267.
- [2] Edward Joyce, "Software Bugs: A Matter of Life and Liability", *Datamation*, May 15, 1987, p. 88-92.
- [3] Janet S. Zeide et al., "Using Expert Systems: The Legal Perspective", *IEEE Expert*, Vol. 1, No. 5, Spring 1987, p. 19-21.
- [4] Richard M. Lucash, "Legal Liability for Malfunction and Misuse of Expert Systems", *SIGCHI Bulletin*, Vol. 18, No. 1, July 1986, p. 35-43.
- [5] Max Stul Oppenheimer, "Liability and Expert Systems", *PC Tech Journal*, Vol. 5, No. 4, April 1987, p. 185-187.

Further Reading

- [1] Raymond Nimmer et al., "Computer Error and User Liability Risk", *Jurimetrics Journal*, Vol. 26, No. 2, Winter 1986, p. 121-137.
- [2] Susan H. Nycum et al., "Artificial Intelligence and Certain Resulting Legal Issues", *The Computer Lawyer*, Vol. 2, No. 5, May 1985, p. 1-10.
- [3] Richard T. Oakes et al., "AI Morality: Machines We Hold Accountable", *Law/Technology*, Vol. 19, No. 4, p. 3-25.

Conference '88



C S C S I '88
Graphics Interface '88
Vision Interface '88

The Convention Centre
Edmonton, Alberta
6 - 10 June 1988

ADVANCE PROGRAM

For General Information:

Wayne A. Davis, *General Chairman*
Department of Computing Science
University of Alberta
Edmonton, Alberta T6G 2H1
Tel: 403-432-3976

TUTORIALS

8:30 a.m. - 5:00 p.m. - Monday - 6 June 1988

#1. Computer Vision - T. Kasvand, Concordia University & NRC; M. Rioux, NRC; P. Boulanger, NRC; A.K.C. Wong, University of Waterloo

This tutorial will cover various aspects of computer vision, particularly those aspects that are related to the theory and practice of 3-D image processing. Topics that will be covered include: data acquisition, differential geometry, surfaces and curves in 3-D space, detection of edges, coherence and similarity, segmentation, 3-D descriptions, recognition of 3-D objects.

#2. The Development of Knowledge-Based Systems - Brian Gaines, University of Calgary

A brief introduction to Artificial Intelligence, and a survey of its accomplishments and current applications will be presented in this tutorial. Topics to be covered include: practical issues in the development of expert and hypermedia systems, potential and limitations of existing technology, tools and techniques for rapid prototyping, and future directions and prospects.

Brian R Gaines is Killam Memorial Research Professor and Director of the Knowledge Science Institute at the University of Calgary. He is editor of the International Journal of Man-Machine Studies, the Computers and People book series and Future Computing Systems.

#3. Computer Animation - Daniel Thalmann, Professor, Département d'informatique et de recherche opérationnelle, Université de Montréal; Nadia Magnenat-Thalmann, MIRALab, Montréal

This tutorial will present theoretical and practical aspects of computer animation. The three main techniques of computer animation will be presented: image key-frame animation, parametric key-frame animation and algorithmic animation. Emphasis is placed on three-dimensional animation of articulated bodies. In particular, the role of synthetic actors is explained in detail using a case study: the film 'Rendez-vous a Montréal'. Task level problems are also explained and an evolution towards a fifth generation animation system is presented.

Daniel Thalmann is a Professor in the Department of Information and Operational Research at the University of Montreal. He has co-produced several computer-generated films, one of which has won several awards.

#4. Object-Oriented Graphics - Peter Wißkirchen, Institute for Applied Information Technology, Gesellschaft für Mathematik und Datenverarbeitung

The aim of the tutorial is to demonstrate how tools for the prototyping of man-machine interfaces provided by object-oriented programming can be used in the field of computer graphics programming, and how they differ from classical approaches. Topics will include: the object-oriented programming paradigm; object-oriented environments (Smalltalk-80, Flavor System), existing graphics classes, window systems and icons; segmenting and naming in object-oriented systems; how inheritance can be used in computer graphics programming; the Mode-View-Controller triad of Smalltalk-80; realization of graphic interfaces using existing basic classes; application frameworks and object-oriented prototypers; and comparison with classical methods, e.g., GKS-oriented programming.

Peter Wißkirchen is a director of the Institute for Applied Information Technology of the GMD West Germany. He coordinates a number of research projects in the areas of computer graphics, man-machine communications and integrated office systems, where a common theme is the attempt to utilize a so-called "Knowledge-based" approach.

TUTORIALS

8:30 a.m. - 5:00 p.m. - Tuesday - 7 June 1988

#5. Graphic Design for Computer Graphics - Aaron Marcus, Aaron Marcus and Associates, Berkeley, CA

This tutorial will provide extensively illustrated lectures on detailed principles, case studies, and guidelines for the use of information-oriented graphic design in user interfaces and hard copy. The goals are to acquaint participants with a wide body of available knowledge, to provide practical advice that is immediately useful, and to highlight potential research topics in visible language programming for user interface design and information graphics. Participants will observe and analyze techniques for making displays and hardcopy more intelligible, functional, aesthetic and marketable.

Aaron Marcus is an internationally recognized authority on graphic design for computer graphics, especially chart, form, document, icon and screen design. He has given tutorials at SIGCHI, Siggraph, NCGA and Micograph (Tokyo) conferences in addition to tutorials at companies in the U.S.A., Canada, Europe, Israel and Japan. He and his staff have designed and critiqued computer graphics presentations, user interfaces, templates and documentation for many national and international companies.

#6. Display of 3-D Biomedical Datasets - Anthony Reynolds, Dynamic Digital Displays, Inc., Philadelphia, PA

Computer tomography (CT), magnetic resonance (MR) and positron emission tomography (PET) provides accurate representations of complete organ systems that can be extracted and displayed to obtain realistic views. This type of presentation is becoming increasingly useful in diagnosing abnormalities or in planning of complex surgical procedures or radiation therapy for cancer patients. Topics to be covered will include: data acquisition; presenting 3-D information; segmentation techniques; object representation; rendering algorithms; image quality considerations; object and image space shading techniques; implementations; commercial systems; futuristic approaches, resampling memories, parallel processors.

R.A. Reynolds is director of clinical applications at Dynamic Digital Displays, Inc. He received his Ph.D. from the University of Pennsylvania in 1985 and since then has written a number of papers on 3-D biomedical datasets, which have appeared in several well-known journals and have been presented at conferences in the U.S. and Europe.

#7. Mobile Robotics - Rodney A. Brooks, MIT AI Laboratory, Cambridge, Massachusetts

This tutorial will survey basic concepts and principles of robotics as well as new trends in advanced robotics research. The structure and capabilities of typical robot systems will be discussed using examples from major research laboratories. Also examined will be fundamental methods of manipulation, including programmed motion, compliant motion, grasping and the use of sensory feedback. The tutorial will include work on model-based programming and automatic planning for mobile robots.

CONFERENCE PROGRAM

Wednesday Morning - 8 June 1988

CSCSI

08:30 Invited Speaker

Finding Proofs, Programs & Plans

W. Bibel, University of British Columbia

10:00 Natural Language

Using Default Logic to Derive Natural Language Presuppositions

R.E. Mercer, University of Western Ontario

Expressing Unrestricted Grammars By Extended DCG

E. Knudsen, SYSLAB, University of Stockholm

Graphics Interface

08:30 Animation

Real-Time 4D Animation on a 3D Graphics Workstation

C.M. Beshers & S.K. Feiner, Columbia University

Techniques for Interactive Manipulation of Articulated Bodies Using Dynamic Analysis

D. Forsey & J. Wilhelms, University of Waterloo

MML: A Language & System for Procedural Modeling & Motion

M. Green & H. Sun, University of Alberta

Joint-Dependent Local Deformations for Hand Animation & Object Grasping

N. Magnenat-Thalmann, et al., MIRALab, Université de Montréal

11:00 Invited Speaker

A System for Conducting Experiments Concerning Human Factors in Interactive Graphics

K. Booth, University of Waterloo

Vision Interface

08:30 Robot Vision

Inspection of the Alignment of Surface Mount Assembly Circuit Board Masks

D. Gauthier et al., McGill University

IRLM: An Imaging System for Object Recognition, Localization & Motion Recovery

J.J. Wu & T.M. Caelli, University of Alberta

Processing Single Range Profiles from a Wrist Mounted Laser Range Finder

C. Archibald, NRC Ottawa

Segmentation of Range Images By Piecewise Approximation with Shape Constraints

J.-M. Beaulieu, Université Laval

10:30 Biomedical Imaging

Geometric Unwarping for Digital Subtraction Mammography

Z. Zhou & R. Gordon, University of Manitoba

Computed Radial-Current Topography of the Brain

Z. Koles et al., University of Alberta

Automated Detection of Breast Tumors

S.M. Lai et al., University of Alberta

CONFERENCE PROGRAM

Wednesday Afternoon - 8 June 1988

CSCSI

13:30 Knowledge Representation

Statistically Founded Degrees of Belief

F. Bacchus, University of Alberta

All the Seats Are Already Taken - A Defense of Probabilistic Logic

R. Aleliunas, Simon Fraser University

Generic Strategies & Representations for Communications Networks Sales

I.A. Ferguson & D.R. Zlatin, Bell Northern Research Ltd.

On Using Modal Structures to Represent Extensions to Epistemic Logics

S. Hamilton & J.P. Delgrande, Simon Fraser University

A solution to the Paradoxes of Confirmation

J.P. Delgrande, Simon Fraser University

16:00 Invited Speaker

Non-Monotonic Reasoning: Is the Answer Harder than the Question?

D. Etherington, AT & T Bell Labs

Graphics Interface

13:30 User Interfaces

Squish: A Graphical Shell for UNIX

T.R. Henry & S.E. Hudson, University of Arizona

A Hypertext Environment for UNIX

P. Prusinkiewicz & J. Hanan, University of Regina

A Graphical Perspective on Robot Workcell Programming

P. Freedman et al., McGill University

15:30 Applications

A Graphical Data Model for CASE & CAD

A.J. McAllister et al., University of New Brunswick

Realistic Presentation of Three-Dimensional Medical Datasets

R.A. Reynolds et al., Dynamic Digital Displays Inc.

Three-Dimensional Display of Medical Images

Y.-W. Tam & W.A. Davis, University of Alberta

Vision Interface

13:30 Invited Speaker

Industrial Machine Vision-Reflecting on the Past Decade

J.L. Mundy, General Electric Company

15:00 Remote Sensing & Computational Vision

Data Fusion & Object Recognition

D. Goodenough, CCRS Ottawa

Digital Terrain Models-Remote Sensing Assignments

M. Lasserre, CCRS Ottawa

Symbolic Knowledge Representation for Remote Sensing

G. Plunkett & D. Goodenough, CCRS Ottawa

On Deriving the 3D Surface of the Earth from SAR

B. Guindon, CCRS Ottawa

Spatial Information Extraction with Pixel Swapping

J. Isaka, CCRS Ottawa

CONFERENCE PROGRAM

Friday Morning - 10 June 1988

CSCSI

08:30 Perception II

Speaker Normalization & Automatic Speech Recognition

Y. Bengio & R. de Mori, McGill University

09:30 Invited Speaker

Recognizing Objects with Curved Surfaces & Moving Parts

D. Lowe, University of British Columbia

11:00 Invited Speaker

Neural Networks, Markov Models & Programming in Automatic Speech Recognition

R. de Mori, McGill University

Graphics Interface

08:30 User Interface Design

Telemac: A Remote-controlled User Interface Server

G. Gagnon & J. Vaucher, Universite de Montreal

Eva: An Event Driven Framework for Building User Interfaces in Smalltalk

J. McAffer & D. Thomas, Carleton University

Constraint Grammars in User Interface Management Systems

B.T. Vander Zanden, Cornell University

10:30 User Interface Design II

Virtual Control Panels

P. Prusinkiewicz & C. Knelsen, University of Regina

A Grid-Based Approach to Automating Display Layout

S. Feiner, Columbia University

Decoupling Pointer & Image Functions of Cursors in Space, Time, & Availability

M.J. Muller, Bell Communications Research

Vision Interface

08:30 Invited Speaker

Space Station-An Application for Computer Vision

K.H. Doetsch, NRC, Ottawa

10:00 Applications I

Unsupervised Segmentation of Textured Images

A. Khotanzad & J.Y. Chen, Southern Methodist University, Dallas

A Representation of Structure from Gradient-Based Motion Algorithms

R. Eagleson, University of Western Ontario

A Global System for Matching 3-D Range Data Objects

N.N. Abdelmalek, NRC Ottawa

Speech Modeling By Correlation Fitting

J. Huang, Northwestern Polytechnical University, Xi'an, PRC

Multiresolution Modeling & Matching Using Salient Features

Z.N. Li & Q. Song, University of Wisconsin-Milwaukee

CONFERENCE PROGRAM

Friday Afternoon - 10 June 1988

CSCSI

13:30 Reasoning II

Context Resolution: A Computational Mechanism for Intelligent Backtracking

J.H. You & Y. Wang, University of Alberta

Search Strategies for Conspiracy Numbers

N. Klingbeil & J. Schaeffer, University of Alberta

Instance-based Prediction of Real-valued Attributes

D. Kibler & D.Aha, University of California, Irvine

Iterative Constructs in Non-linear Precedence Planners

S. Steel, University of Essex

15:30 Applications

Exploiting Fine-grained Parallelism in Production Systems

B.T. Smith & D. Middleton, NASA

Qualitative Modeling: Application of a Mechanism for Interpreting Graphical Data

S. McIlraith, Alberta Research Council

Graphics Interface

13:30 Invited Speaker

Computer-Assisted Composition & Direction of Dance, Film & Video

T. Calvert, Simon Fraser University

15:00 Algorithms

The Application of Area Antialiasing on Raster Image Displays

Y.-T. Chen & P. D. Fisher, Michigan State University

Attractors & Repellers of Koch Curves

P. Prusinkiewicz & G. Sandness, University of Regina

A General Algorithm for 3-D Shape Interpolation in a Facet-Based Representation

T. M. Hong et al., MIRALab IRO

Theoretical Results on the Priority Approach to Hidden-Surface Removal

P. Eged, McGill University

Vision Interface

13:30 Applications II

Parallel Algorithms for Image Object Labeling

E. Leung & X. Li, University of Alberta

Syntactic Pattern Analysis in Seismic Exploration

L. Le & E. Nyland, University of Alberta

Label Relaxation Technique for Stable Estimation of a Topographic Primal Sketch

P. Boulanger, NRC Ottawa

A Low Cost Robot Vision System: Successes & Shortcomings

E.J. Trost & S. Onyshko, Unisys & University of Manitoba

Feasibility of a Global Description of 3-D Objects in Range Data for Recognition

S.H.Y. Hung, NRC Ottawa

16:00 Invited Speaker

Integrating Methodologies in Image Analysis

T. Pavlidis & Y.T. Liow, SUNY, Stony Brook, New York

Conference '88 Registration

Name _____

Phone Number _____ Affiliation _____

Address _____

Society Membership CSCSI CMCCS CIPPRS Membership # _____

Which set of proceedings would you like (one free per registrant)? CSCSI '88 GI '88 VI '88

Which tutorials would you like to attend (use tutorial numbers)?

FEES - circle applicable items

	Before April 1 Member/Non-Member	After April 1 Member/Non-Member	Student Member/Non-Member
Conference Only			
3 days	\$215/255	\$265/305	\$95/130
Tutorials Only			
1 day	155/185	185/215	70/90
2 days	255/305	305/355	110/150
Conference & Tutorials			
5 days	380/440	440/495	110/140
Proceedings - extra	15/30	25/35	20/30
Film Show Only		10	
BBQ Tickets - extra		35	

Total Enclosed \$ _____

Payment to Conference '88 by cheque or money order must accompany advance registrations. Send conference and hotel registration to

**Conference '88 Registration
Department of Computing Science
University of Alberta
Edmonton, Alberta
Canada T6G 2H1**

3 day Conference registration includes: 1 proceedings, reception, barbeque, film show

Student registration includes: 1 proceedings, reception, film show, barbeque. Students must be full-time and provide valid ID at registration.

Tutorials registration includes: notes for registered session, luncheon

All prices are expressed in Canadian Dollars. In U.S. Dollars, prices are 20% less.

Refunds will be given only if notification is made in writing by 23 May 1988, and are subject to a \$20 service charge.

HOTEL REGISTRATION FORM

Reservations must be made using this housing form and mailed no later than 1 May 1988. All reservations are to be accompanied by a deposit equal to one night's stay. The deposit may be made either by cheque, money order or credit card payable to: **Edmonton Convention & Tourism Authority**. U.S. or Canadian funds may be used. *No phone reservations accepted.* Confirmations are mailed out directly from the hotel.

Mail this form, along with your conference registration to **Conference '88 Registration** (address on reverse). To change your arrival/departure date or to cancel your reservation, contact

Edmonton Convention & Tourism Authority
#104, 9797 Jasper Avenue
Edmonton, Alberta,
Canada T5J 1N9 (403)426-4715

before 27 May 1988. After that date, contact your hotel directly. We recommend making all changes in writing. Reduced conference rates are available only through the use of this form.

Name _____ SharingWith _____
Address _____ Phone Number _____

City _____ Prov/State _____ Postal/Zip Code _____

HOTEL PREFERENCE

Please place the numerals 1, 2 and 3 in the spaces below to indicate hotel choice.

Note that all requests are processed on a first-come, first-served basis.

Chateau Lacombe	Four Seasons	Westin
_____ \$65.00 single/dbl	_____ \$85.00 single	_____ \$90.00 single/dbl
	_____ \$90.00 double	

Arrival: Date _____ Time? am pm Departure: Date _____ Time? am pm

Single room? _____ Double room? _____ Non-smoking Room? _____

Card: Am. Ex. MasterCard Visa Number _____ Expiry Date _____

Signature _____

Student accommodation is available from Conference Center,
#4 Lister Hall, University of Alberta, Edmonton T6G 2E1
(403) 432-7200

Hotel reservations will not be made unless accompanied by a completed conference registration.

MAPLES and EVES

at I.P. Sharp Associates (Ottawa)

by Graeme Jones and Bill Pase

MAPLES et EVES à I.P. Sharp Associates (Ottawa)

RÉSUMÉ: Le Groupe de recherche et technologie d'Ottawa d'I.P. Sharp Associates Ltée. poursuit présentement deux projets associés à l'intelligence artificielle. MAPLES est un ensemble d'outils logiciels qui peuvent être utilisés par les cognitivistes pour développer des systèmes experts en APL. EVES est un système de vérification de programme en cours de développement pour le Ministère de la défense nationale Canada et la U.S. Navy.

The Ottawa Research and Technology Group of I.P. Sharp Associates Limited is currently involved in two projects related to artificial intelligence. The first, MAPLES, is a tool for developing APL-based expert systems. The second, EVES, is a program verification system.

The team consists of seven researchers, one technical writer, one administrative staff member, and a project manager. The hardware resources available for the projects include an IBM PC/XT, an AT&T 3B1, two VAX 750s, and four Symbolics Lisp Machines. The software resources include SAX (Sharp APL for Unix), STSC APL*PLUS, VAX/VMS, Berkeley Unix, and the Genera system. The two VAXen and the four Symbolics reside on an ethernet. One VAX is also connected to the ARPANET.

MAPLES

MAPLES (Mixed paradigm APL-based Expert System toolkit) is a set of software tools, which can be used by knowledge engineers to develop APL-based expert systems. It is intended to achieve the following goals:

- increase the availability of expert systems technology to APL users
- determine the strengths and weaknesses of APL as a language in which to implement expert systems
- provide the Communications Security Establishment (CSE) with a software tool for building risk analysis expert systems.

A first set of requirements of the software tools has been jointly established by I.P. Sharp Associates Limited and CSE. The work plan is to develop product versions of MAPLES for two APL systems: SAX (Sharp APL for UNIX) running on an AT&T 3B1 computer, and STSC APL*PLUS running on an IBM PC/XT. A MAPLES prototype for the SAX environment will be demonstrated by I.P. Sharp Associates Limited at a SigAPL workshop (Syracuse, New York) in August.

MAPLES consists of a knowledge-base management facility, STAPLE (Structured TURING and APL Environment), and a set of expert system utilities. The knowledge-base management facility will enable knowledge engineers to construct and maintain databases using the following paradigms for

Graeme Jones and Bill Pase are researchers at I.P. Sharp Associates, Ottawa.

representing knowledge: frames, production rules, predicate calculus, and the STAPLE programming language. The latter will be used to define value types of frame slots. Either STAPLE or APL may be used to specify the routines (inference methods, procedural attachments, and end-user interfaces) of MAPLES expert systems.

STAPLE is an environment for developing programs in the STAPLE Programming Language, which combines TURING (from the University of Toronto) with the expression language of APL. The environment will include a compiler for translating STAPLE programs into either Sharp APL or STSC APL*PLUS, depending on the APL system being used. The set of expert system utilities will include rule-chaining inference methods and standard abstract data types. The STAPLE compiler will be written in UNIX C (using Lex and Yacc) for the SAX environment and translated into Turbo C for the PC.

The following work on the project has been completed:

- the formal software requirements of MAPLES
- a formal methodology for representing and acquiring knowledge (to be implemented in MAPLES)
- the design of the STAPLE language.

EVES

EVES (Environment for Verifying and Evaluating Software) is a program verification system being developed for the Canadian Department of National Defence and the United States Navy. There are two major streams to the development of EVES: the design of a specification and implementation language (with the supporting mathematics), called Verdi; and the implementation of a theorem prover, called NEVER.

The first phase of the research and development has been completed. This has resulted in the m-EVES system which consists of the language m-Verdi, an m-Verdi compiler, and the m-NEVER theorem prover.

The m-Verdi language allows for the writing of formal specifications and the implementation of imperative programs. The formal semantics for m-Verdi is described using a form of Denotational Semantics. The logic is non-standard in that it extends the many-sorted Predicate Calculus so that new symbols can be added to a theory (using m-Verdi declarations), and includes proof obligations requiring the "conservative extension" of

theories. The logic requires that programs be proved to terminate and that the specification of program functionality be of the well-understood Floyd-Hoare style.

A compiler for m-Verdi has been written and runs on VAX computers under the VMS or UNIX operating systems. The compiler optimizes the generated code and is easily retargetable to other machines and operating systems.

The m-NEVER (Not the EVES Rewriter) theorem prover supports the interactive development of proofs and also has powerful automatic capabilities. These capabilities include the detection of propositional tautologies, the handling of equality and integer relations, the application of conditional rewrite rules, the application of forward chaining rules, the heuristic expansion of defined functions, the automatic instantiation of quantified variables, and the Boyer-Moore style of automatic induction.

The theorem prover development has liberally borrowed ideas from work by Bledsoe, Boyer and Moore, Nelson and Oppen, and the work embodied in Affirm. A detailed presentation of m-NEVER is available.

Future work on the theorem prover will be directed at the development of a proven "proof checker" for m-EVES proofs. Such a checker will take m-EVES proofs and ensure that each inference is sound.

The entire m-EVES system, except for the m-Verdi compiler, currently resides on a Symbolics Lisp Machine. Interaction with m-EVES may occur through either an editor interface or a command processor.

The development of m-EVES was completed in mid-November, 1987. Future work (resulting in the development of EVES) will be directed at increasing the expressibility of the logic and m-Verdi (resulting in Verdi), the writing of a compiler for Verdi, the continuing evolution of the theorem prover (resulting in NEVER), the continuing evolution of the interface, the porting of the system to other hardware bases, and the continued application of the system to various examples.

Personnel

Researchers: Dan Craigen, Graeme Jones, Sentot Kromodimoeljo, Irwin Meisels, Bill Pase, Mark Saaltink (and, until recently, Andy Neilson). Technical writer: Karen Summerskill. Administrative staff: Denine Wrixon. Project manager: David Bonyun.

Intelligent Tutoring Systems: the HERON project

by Claude Frasson

Systèmes d'Enseignement Intelligemment Assisté par Ordinateur: le projet HERON

RÉSUMÉ: Les recherches en enseignement intelligemment assisté par ordinateur (EIAO) résultent d'une convergence de deux communautés distinctes: l'Intelligence Artificielle et l'Éducation. L'éducation représente un domaine très large d'expérimentation pour l'IA étant donné qu'elle touche à des problèmes de raisonnement et d'acquisition de connaissances. Ces problèmes demeurent complexes car ils sont reliés au comportement humain, domaine qui n'est régit par aucune règle mathématique mais plutôt par des "théories" ou "méthodes" d'enseignement fort controversées. L'utilisation de techniques d'IA en Éducation peut avoir un double impact: elle peut contribuer à une meilleure compréhension des stratégies efficaces d'enseignement et d'apprentissage, et elle contribue aux fondements des connaissances (acquisition de connaissances, raisonnement, déduction...). Considérant que des environnements avancés d'apprentissage doivent bénéficier d'une variété d'expertise telles que des connaissances en bases de données, génie logiciel, apprentissage, interaction multimédia, psychologie et démonstration de preuves, nous avons réuni une équipe multidisciplinaire à l'université de Montréal au sein du projet HERON. Le projet regroupe environ 12 chercheurs et 11 étudiants. La moitié des chercheurs proviennent de l'université de Montréal, les autres de l'université du Québec à Chicoutimi (UQAC), de l'université du Québec à Montréal (UQAM), de l'université Bishop's du Collège Militaire Royal de Kingston, et du Open System Group de Victoria. Les principaux axes de recherche de HERON concernent: la théorie de la connaissance, les stratégies d'enseignement, les bases de connaissances distribuées et l'utilisation de techniques d'IA pour le contrôle des médias. Le groupe utilise essentiellement des postes de travail SUN, reliés sous ethernet, de forte capacité de mémoire (16-20 Meg) et opérant sous le système ART (Automated Reasoning Tool). Il utilise aussi deux VAX/780 sous Unix, opérant avec INGRES et plusieurs langages de développement. Les expériences sur les médias se font sur un AMIGA 1000, à écran tactile, utilisant vidéodisque, caméra, digitaliseur d'images, magnétoscope et synthétiseur de son.

Since the introduction of Computer Based Instruction in Education in the 1950s, there have been progressive improvements in the feedback given to the student and in the individualization of interaction. Intelligent Tutoring Systems (ITS) effectively began to appear in the 1980's,

Dr. Claude Frasson is a professor in the Dépt. d'informatique et de recherche opérationnelle, Université de Montréal. He is also the conference general chairperson of the upcoming International Conference on Intelligent Tutoring Systems.

as new tools and formal approaches emerged in artificial intelligence, particularly in diagnosis, knowledge representation and inference mechanisms.

Research in Intelligent Tutoring Systems is a convergence of work resulting from two distinct communities: artificial intelligence and education. Education represents a very large domain of application for AI, encompassing problems of both reasoning and knowledge acquisition. The domain of education

however, is not supported by mathematical rules but, rather, by different and much debated "theories" or "methods" of teaching and learning. Furthermore, there are no tools to quantitatively measure the efficiency of each theory and the learning environment is rapidly evolving in terms of new technology, growing flow of information, and so forth. These three factors enhance the complexity of applying AI to the domain of education. The use of AI techniques in education however, can contribute to a better understanding of efficient teaching and learning strategies as well as to the foundations of knowledge (knowledge acquisition, reasoning, and deduction.)

Considering the fact that an advanced learning environment can benefit from various areas of expertise such as databases, knowledge engineering, learning, multi-media interaction, psychology and proof theory, we have formed a multidisciplinary team at the Université de Montréal, under a project called HERON. The project team consists of approximately twelve researchers and eleven students. Half of the researchers are from the Université de Montréal, the others being from Université du Québec à Chicoutimi (UQAC), Université du Québec à Montréal (UQAM), Bishop's University, Royal Military College of Kingston and the Open Systems Group from Victoria. In the following we will discuss the main research projects in HERON.

Theory of Knowledge

The problem of knowledge representation in Intelligent Tutoring Systems is a broad one due to the variety of types of knowledge involved. For instance, there is knowledge on particular subject matter to be transmitted to the student, knowledge acquired by the student, the concepts of partial or total understanding, and the evolution of knowledge along the time and degree of abstraction of the subject matter. Several sub-projects related to this problem are under way; these are discussed below.

A Knowledge Model of the Curriculum

This model includes the subject matter, the student model and student performance. The project involves Gilles Gauthier (UQAM), Claude Frasson (Université de Montréal) and PhD student Gilles Imbeau. An implementation of a prototype in office automation (a course on word processing) is under way with Samy Bengio, Aziz Ouazzani and Jean Luc Landry.

A Model of Reasoning

This is a long term project which is attempting to clarify and structure the principle of knowledge acquisition in order to tackle the different forms of reasoning (induction, association, deduction, analogy). The research team consists of Marc Kaltenbach (Bishop's University), Claude Frasson, Jan Gecsei (Université de Montréal) and Gilles Imbeau. Preliminary work done by Marc Kaltenbach included building a system for teaching, step by step, proofs of mathematical theorems.

Student Modelling

An important component of Intelligent Tutoring Systems is a representation of the knowledge of the student and the capability of the system to formulate a diagnosis of his/her cognitive state. This includes not only psychological aspects but also a variety of

performance measures. PhD student Guy Mineau and Soudougou Konate are investigating how to capture and maintain a student's profile and what performance parameters should be considered.

Teaching Strategies

According to the broad literature on the subject and the variety of theories available, no one seems to be able to explain how to teach well. Several ongoing projects aim to clarify and evaluate teaching strategies.

Harold Stolovitch (Sciences of Education, Université de Montréal) is experimenting with several models of learning in various situations. The objective is to detect through these developments the characteristics which lead to useful and efficient systems. A formative evaluation is also realized using games and simulations.

Andre Ouellet and Richard Tremblay (UQAC) are focusing on the development of a rule-based system to validate a model composed of four inter-related subsystems: the goals of the subject matter, the student's knowledge, the microworld, and the evaluation of performances. The model will be able to detect ambiguities between different learning rules through a sequence of dialogues between the subsystems.

Andre St. Gelais, Gilles Imbeau and Claude Frasson are designing the specifications of an intelligent planner that will be able to generate plans, execute a plan and revise it if necessary. In the next step, the planner will be integrated with the student model and the curriculum in a prototype in office automation.

The use of standard Computer Assisted Instruction Systems (CAI) would be more interesting if "intelligent" capabilities could be introduced. Samy Bengio is investigating methods of combining CAI systems with Intelligent Tutoring Systems. The execution of teaching tasks is devoted to the CAI component while reasoning, evaluation and the planning of the appropriate sequence of tasks are all handled by the ITS.

Roger Hart (Open Systems Group, Victoria) is working on Knowledge Based Learning or "applied systemology" as a distinct subset of Intelligent Tutoring Systems, which seems to offer some interesting solutions to many of the problems faced by educators in the 1990's. The ideas derive from the works of William Morris and Gordon Pask as well as recent approaches in AI.

An ongoing project involving teaching strategies is directed towards the teaching of "strategies." Jean Fugere (Royal Military College of Kingston) and Rejean Carrier are currently designing and building an interactive Intelligent Tutoring System able to explain military tactics and advise officers involved in a simulated battlefield.

Distributed Knowledge Bases

The distribution of knowledge between several fields of expertise can have an impact on the architecture of Intelligent Tutoring Systems. Two projects are directed at that area; these are described below.

Cooperation Protocols

Gregor Bochmann (Université de Montréal) is involved in research on cooperation protocols between the components of a distributed knowledge base. The components may be of the same type of knowledge or not.

Composition of Knowledge Bases

PhD student Mohammed Er-radi is designing and building an intelligent system able to select and extract knowledge from a distributed knowledge base, and compose all the elements needed for a specific course. This includes presentations, exercises, video-demonstrations, etc.

AI Techniques for Media Control

Videodisc technology offers a remarkable combination of frame capacity and access time which represents a useful tool in education. The potential of this technology can be enhanced by employing AI techniques for interaction with the media. Two projects are underway.

Jan Gecsei, Daniel Martin and Alain Pilon are experimenting with browsing techniques for video databases using the "fish-eye" principle. A fish-eye view consists of a focus image (still or moving), in full video resolution on the entire screen, with still images of the neighbourhood displayed as miniatures around the edges of the screen. Using a touch screen, browsing is done simply by pointing to a miniature or to an area in the focus (in order to obtain an enlarged version or descriptive information). Pointing to the center of the focus triggers the movement (if the document is a

sequence). The choice of the set of miniatures (all related to a common topic) can be done using AI techniques.

PhD student Paul Girard is exploring ways of integrating and synchronizing sound and real-time animation in a teaching interaction. Using an images database and a sound synthesizer, comments, oral presentations and explanations can be automatically generated. With such a system, the student has the feeling of being online with his teacher. The researchers have been experimenting with a synthesizer able to speak various languages.

Computing environment

The HERON group is equipped with a number of SUN workstations fully networked by Ethernet. Three of these workstations are running ART (Automated Reasoning Tool) with sixteen to twenty megabytes internal memory. ART is being used for the development of all of the prototypes. The group also has the use of two VAX 11/780s running Unix, Ingres and a variety of development language software. All the HERON members are linked by electronic mail. The experiments on media are done on a Amiga 1000 (with touch screen), using a pioneer LD-V6000A videodisc player, cameras, image digitizer, VHS and a synthesizer.

Artificial Intelligence Research and Applications at the Alberta Research Council

Edited by Sheila McIlraith

Recherches et applications en intelligence artificielle
au Conseil de Recherche de l'Alberta

RÉSUMÉ: Le groupe des technologies d'information du département des technologies de pointe du Conseil de Recherche de l'Alberta s'occupe d'IA, de calcul réparti, et d'autres domaines informatiques. L'expérience et le savoir de son personnel constitue son plus grand atout. Le groupe comprend deux doctorats et sept maîtrises en informatique ainsi que six baccalauréats allant de l'informatique à la géologie et la psychologie. On dispose d'un matériel de pointe dont sept Symbolics, cinq Suns, plusieurs MacIntosh et IBM PC, un MicroVax de DEC et un ordinateur Prime. Le logiciel disponible comprend ART, Knowledge Craft/Language Craft, Goldworks, Nexpert et Object Nexpert. On utilise aussi le PROLOG de Quintus, plusieurs LISPs, Smalltalk et C. Le travail du groupe passe à l'industrie par quatre voies: 1) le programme de recherche interne, 2) des contrats industriels, 3) le programme JRV d'entreprise de recherche conjointe et 4) les contacts extérieurs. L'article donne plus de détails. Jusqu'à date la recherche a porté sur la planification de processus, l'IA répartie, l'IA en éducation et les systèmes experts.

The Alberta Research Council (ARC), a provincial Crown Corporation, was founded in the early 1920s. It is the largest and oldest Canadian provincial research organization, having established a world class research reputation in a number of areas including oil sands and biotechnology. The ARC mandate is to advance the economy of the province by promoting technology development, performing applied research and providing expert advice, technical information and scientific

This article is a compilation of material from several papers and internal reports, with structuring and additional text written by the editor. Sheila McIlraith is an associate researcher in artificial intelligence at the Alberta Research Council.

infrastructure that is responsive to the needs of the private sector and supports activities in the public sector. Hence, ARC's activities tend towards applied research and technology transfer.

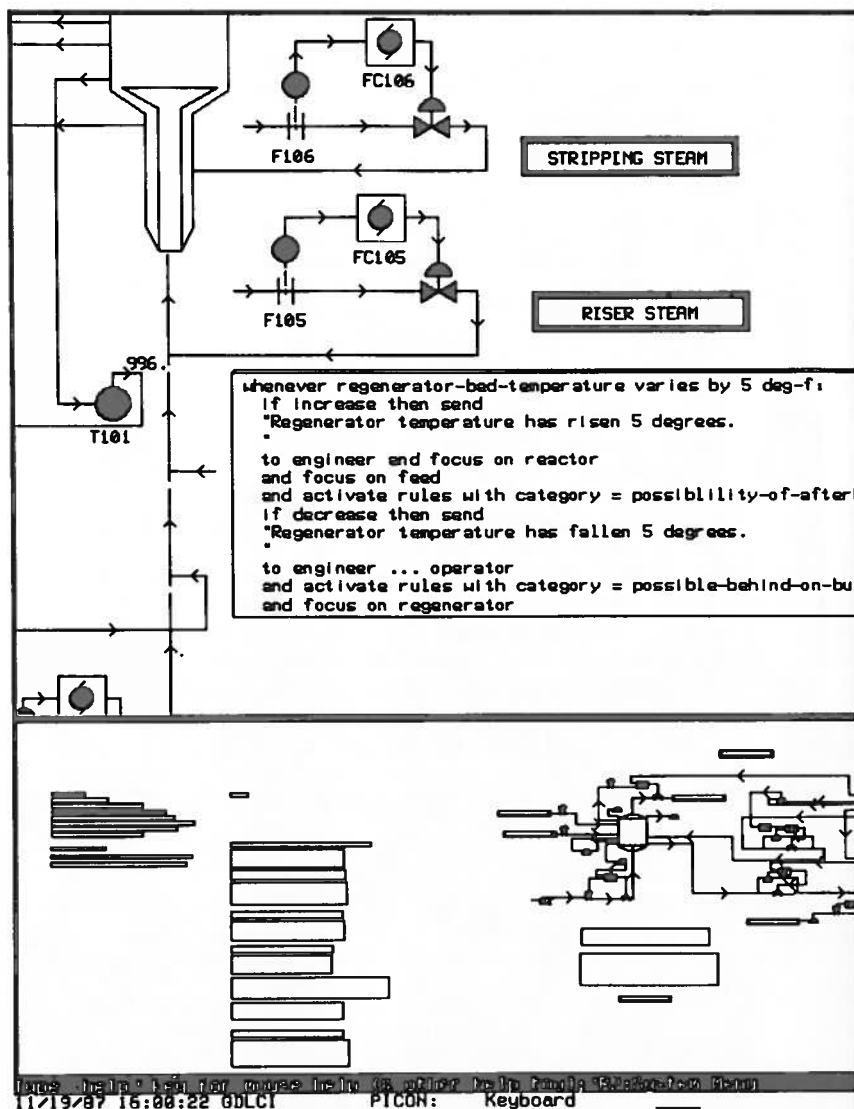
The Advanced Technologies Department, located in Calgary, was formed in 1985 with a mandate to lead in the diversification of the Alberta economy through technological investigation and industrial liaison in information and materials handling. This led to the development of two technical groups: Information Technologies and Automation Technologies. The Information Technologies group focuses on research and development in artificial intelligence, distributed computing, and other areas of computer science.

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 an application system. It is developed from expert systems and process control.

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

GigaMos Systems Inc.

a leader in the application of artificial intelligence.

PICON


Load
Save
Merge

Run
Reset

Visit
Retrieve
Print

Add Icon
Add Statement
Add Table

irny



*rule-driven, definition-based
experience in artificial intelligence*

LISP Canada Inc.

1000 Roche Boulevard
Vaudreuil, Québec J7V 6B3

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

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, DATA-BASE 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.

Automation Technologies performs research and development in robotics, signal processing, system integration, and other related engineering activities.

Information Technologies, the centre for AI activities at the Alberta Research Council, has an experienced and knowledgeable research and programming staff, including two Ph.D.s and seven M.Sc.s in computer science, and six B.Sc.s in disciplines including computer science, geology, and psychology. State of the art computing facilities include seven Symbolics workstations, five Sun workstations, Macintosh and IBM personal computers, a DEC Micro Vax II, and a Prime computer. A variety of software resources are available to staff for application development. ART, Knowledge Craft/Language Craft, Goldworks, Nexpert and Object Nexpert are some of the available tools. In addition, a number of programming languages (including Quintus PROLOG, a variety of LISPs, Smalltalk, and C) are provided on a variety of hardware systems.

Information Technologies' major activity is in the area of research and technology transfer as described below. In addition, ARC is involved in a number of AI related initiatives. ARC is a member of, and is represented on the board of directors and technical advisory committee for PRECARN, a national research consortium in artificial intelligence and robotics. Additionally, the Alberta Research Council has signed a research and development agreement with International Artificial Intelligence to be supported in our research by the provision of 5 Symbolics workstations. The Council is also actively involved in the Canadian Space Program through the presentation of a statement of interest, and through the submission of joint RFP responses with industry. Finally, members of Information Technologies enjoy publishing and editing a well known Canadian artificial intelligence magazine, when not otherwise occupied.

Research and Technology Transfer

Information Technologies staff perform research and technology transfer through four channels: (1) the internal core research program, (2) direct contracts with industry, (3) the ARC joint research venture (JRV) program, and (4) external research associations.

In addition, the Alberta Research Council actively supports technology transfer to the industrial sector, especially that of Alberta. In response to current needs, industrial liaison has been directed largely at the area of expert systems. Information Technologies provides instructional video tapes, 1-day expert systems workshops, free expert systems clinics, self-study programs, public seminars, demonstrations, discussions with interested parties and potential clients, and other related activities. Also available to interested individuals is the Alberta Research Council's extensive collection of journals and books.

Listed below are a number of the research and development projects underway within Information Technologies. The names of ARC participants are listed below each project. In some cases, particularly with the larger contracts and Joint Research Ventures, there has been significant participation by staff members from the contracting companies.

1. Core Research Program

The core research program, supported internally by ARC, was developed to perform research in Information

Technologies, in response to the current and anticipated future needs of Alberta industry. Some of the projects listed are in direct response to inquiries from industries, others have generated industry interest in the nature of problem specification or financial assistance.

Process Planner II

The Process Planner II, the largest individual research project in Information Technologies, is a multi-discipline AI project concerned with the specification, planning and execution of changes in physical systems involving time. It seeks to apply current research results from the areas of temporal reasoning and planning, knowledge representation of objects and processes, and natural language understanding towards building a knowledge-based system for planning processes. In a typical manufacturing plant, after a product is designed, a process engineer determines the activities necessary for manufacturing the product. This usually involves adapting the currently known processes to the needs of the new product. The Process Planner II is aimed at providing these engineers with an intelligent tool for use in process planning. After an initial training period where II is customized to a specific environment, the system will be able to generate process plans for various products manufactured in the plant. An engineer specifies components and the processes operating on them, and gives a description of an end-product. II then generates the specific sequence and duration of the processes. Its output will be a synchronized instruction stream for the appropriate effectors.

The system is comprised of two main parts: a representation of the environment, and a planner. The environment representation consists of four components: objects, agents, tools, and processes. For a given environment, it is then necessary to plan the application of the processes using the tools and processes, applied by agents to the input objects in order to produce the desired end-object. The planner is responsible for determining the sequence of activities necessary for assembling a product from its component parts; the planning paradigm is based on least commitment and constraint satisfaction techniques. Finally, the output of the planner must be translated into the appropriate instructions given the agents, tools and processes that will be used. This instruction stream can be single, or parallel for multiple agents, and at various levels of complexity, depending on the capability of the agents, which may be simple tools, robots, or humans.

Conceptualization and design of the system are completed [1]. At the root of II is James Allen's temporal logic representation. A temporal reasoner and rudimentary planner have been implemented in Knowledge Craft on a Symbolics 3620. Work continues on improving the planner and implementing the knowledge representation scheme for the environment.

Participants: Ken Gamble, Roy Masrani, Sheila McIlraith

Participant Systems

The Participant Systems project is concerned with research and development in distributed computer systems that facilitate the interaction of multiple agents, both human, and machine-based. This research involves work in the human-computer interface, computer graphics, computer communications and artificial intelligence, with background knowledge of cognitive

and social psychology.

A Participant System [2] is a computer system that facilitates the simultaneous interaction of several persons working together, possibly over several physical locations, on a shared complex task. To do so, it must support communications, multiple views, and common data, action and cognitive space. Such a system must coordinate access to a common problem representation, and contain sufficient knowledge and expertise in the problem domain to integrate the activity of the users, and even to participate as one of the experts. This represents a new paradigm for computing, a departure from the traditional one-person-to-one-virtual-machine model, to a model paradigm that more directly reflects the nature of human problem solving in the group situation.

Two systems have been developed as testbeds for the concepts involved in Participant Systems. Cantata is a Macintosh-based system that supports group communication and use of computer resources. It is intended to be a vehicle for research into group interaction protocols using computer-based media.

The other program under development on the Macintosh is the Participant Construct System (PCS), a knowledge acquisition tool which operates in either a local or network environment. It is based on the notion that knowledge resides in the criteria upon which we make distinctions among similar entities. PCS can be used for a wide range of knowledge-based applications from obtaining rules for expert systems, to obtaining consensus among groups, to identifying selection criteria where choice is involved. It also has potential for application in the computer-assisted design process.

Current research in Participant Systems is in the development of a quantitative measure for the effectiveness of protocol mechanisms in the performance of collaborative tasks by a group. Apple, Cupertino, is supporting this project with a grant for some MAC IIs, in addition to graduate student support.

Participants: Ernie Chang, Tony Copping, Richard Kasperski, Patricia FitzGerald

ISIM

ISIM, an Intelligent SIMulation system, aims to integrate reasoning and representation techniques developed in artificial intelligence with discrete-event simulation. The project investigates several avenues for integrating AI and simulation: (1) a goal directed expert system that determines the appropriate simulation parameters for achieving pre-stated objectives, and (2) an exploratory system that experiments with the simulation model to find pertinent relationships between components of the model.

In order to achieve these goals, a simulation specification environment has been created by enhancing Knowledge Craft with a Simula/DEMOS-like representation methodology. This environment captures the rich knowledge representation techniques of Knowledge Craft, while taking advantage of past research in formalizing the Simula/DEMOS framework. The ISIM specification environment is augmented by a number of tools. Traditional quantitative analysis techniques such as response surface methodology (RSM) have been implemented in ISIM to assist the expert system in determining appropriate simulation parameters to achieve predefined simulation objectives. In addition, both a qualitative dependency analysis, and a quantitative factor analysis may be performed to give the

user some insight into the relationship between components of the specified simulation model, and system responses. The expert system may use these tools to run an optimal set of experiments for analysis of the specified model. Work on ISIM continues with attention placed on enhancing the expert system component, integrating the existing tools, and improving the user interface.

Participants: Roy Masrani, Sheila McIlraith, Brendan Mumey, Margaret Mendoza

2. Industrial Contracts

Another means by which Information Technologies performs research and technology transfer is through industrial contracts with clients. The Advanced Technologies Department conducts feasibility studies and fee-for-service contracts which enable companies to work with skilled professionals to develop applications that will advance their competitive position. Listed below are some examples of work performed to date. Some projects have been omitted, in other cases, technical details are not described in order to comply with the desires of ARC clients. Some of the work performed has led to interesting applied research results.

An Expert System for Well Test Interpretation

The purpose of this contract was to determine the feasibility of developing an expert system for well test interpretation. Feasibility was established and a multi-expert modified blackboard architecture was proposed. As part of the contract, a prototype qualitative modeling [4] mechanism for interpreting graphical data was designed and implemented. It models and interprets graphical data by identifying characteristic curve shapes that indicate physical phenomena. This system was implemented in ART on a Symbolics 3670.

Participants: Sheila McIlraith

An Expert System for Emergency Response Management

Information Technologies, Alberta Public Safety Services and Emergency Preparedness Canada combined expertise to develop Hermes: An Expert System for Emergency Response Management [5]. Hermes, a prototype system to date, combines knowledge of the properties of dangerous goods, with expertise in the management of chemical spills to provide advice on the management of emergency situations involving dangerous goods. The system, originally implemented in ART on a Symbolics 3620 includes both an expert mode and a tutorial mode. The expert mode accepts an incomplete graphical and textual description of an accident situation, assesses the danger, and suggests corrective actions. Additional information regarding the accident or corrective actions that have been implemented may be input at any time. Hermes will immediately update the assessment of the situation with the assistance of ART's truth maintenance system.

The tutorial mode of Hermes provides a set of lessons for the user. Each lesson provides an emergency scenario and a set of lesson objectives, or skills the student is to improve. Using the Hermes interface, the student must assess the situation and select corrective actions. The student may then be rated on his/her response and receive an explanation of where marks were gained or lost. Hermes is currently being ported to Goldworks on

an IBM-AT. Further development is planned.

Participants: Margaret Mendoza, Lynn Sutherland, Roy Masrani, Greg Sidebottom, Ernie Chang

A Knowledge-based System for Character/Word Recognition

Three consecutive contracts were performed for DataSpan Technologies Ltd., a Calgary-based firm, to augment existing raster-vector conversion software, and vector editing facilities with an *intelligent* orientation independent handwritten character/word recognition system. Techniques from artificial intelligence, syntactic pattern recognition and from work in connectionist systems [6] were investigated during design of the system. Prototypes were initially developed to establish the degree to which various techniques complied with system requirements. A larger system was developed and integrated with client software to perform orientation independent handwritten number recognition. The new digit classification system was enhanced by integration of an AI component that used local contextual knowledge, the known format for the representation of numbers, and expectation of the values of numbers in a particular document, to resolve digit ambiguities and to concatenate adjacent digits to form numbers. The prototypes were developed in a variety of environments; the final system was implemented in C on a Sun. The client company participated actively in development of the system. Work in related areas is scheduled to continue.

Participants: Sheila McIlraith, Tony Copping, Yiu-Wing Tam, Farran Sloan

An Expert System for Configuration of Telephone Services

Information Technologies has performed some work for Alberta Government Telephones (AGT) to design and prototype an expert system for the configuration of telephone services. The expert system, developed in Goldworks on an IBM-AT will assist sales people in delivering telephone services to customers. In addition it will generate forms and related database data for such tasks as quotations, programming forms and sales forms. Related systems are planned for the future.

Participants: Ken Gamble, Darcy Grant

Expert System Augmentation of an Authoring Language

Softwords, a Victoria-based firm, contracted ARC to integrate artificial intelligence techniques into their computer based training software. The task was to enhance NATAL, a procedural authoring language, by interfacing it with ARC's expert system shell, PORTAL on an IBM-AT. The resulting authoring environment allows the expert system shell to represent what is to be taught using rules and deep model reasoning, while the procedural language within NATAL represents how the subject material is to be taught.

Participants: Ringo Ling, Ken Gamble, Marlene Jones, Ernie Chang

Enhanced Ice Image Recognition

Information Technologies has just commenced a contract for Intera Technologies Ltd., a Calgary-based firm. The project is intended to create a prototype software system that will enhance the ability of human

operators to recognize distinct types of ice formations in images generated by Intera's Synthetic Aperture Radar (SAR) system. Information Technologies is to provide advice and assistance in knowledge acquisition, knowledge representation, computer assisted instruction, and human-computer interface design. The ARC knowledge acquisition tool Participant Construct System (PCS) is currently being used for knowledge elicitation.

Participants: Ruby Loo, Ernie Chang

3. The Joint Research Venture Program

The Joint Research Venture (JRV) Program is an initiative of the Alberta Research Council to assist Alberta companies in developing chosen technologies which are relevant to the companies' position in the marketplace. The JRV program focuses on the development of innovative products or processes. The fundamental technology required must be available and there should be a recognized market for the product or process. JRV projects are generally 1-3 years in duration and tend to have a major research component. Joint participation from both ARC and the contracting company is required by definition. Some of the more recent JRVs are listed below.

An Expert System for Well Log Analysis

In January of 1988, D&S Knowledge Systems Inc., a new Calgary firm spawned by the success of the JRV, announced the successful completion of a practical knowledge-based software system which has been used as an "assistant" in the interpretation of well logs. The expert system for well log analysis has been integrated with D&S Petroleum's LOGMATE, a fourth-generation log analysis workstation. Together they provide graphics, database, data communications, arithmetic, and report generation facilities. Given some well log data, the expert system will either automatically provide an analysis of the data, or will suggest data specific analysis methods for the user to invoke using the system. The organization and representation of analysis methods in order to reason about their behaviour and to generate algorithms for new functions is an interesting coupling of symbolic and numeric computing [7]. A full explanation of the suggested analysis methods is also available from the system. The fourth-generation workstation and expert system are open and data-driven; experts may modify the knowledge base to reflect individual analysis preferences. The system is implemented in C and runs on an IBM-AT or compatible, and on UNIX-based microcomputers. The methodology incorporated in the development of this expert system may be used in other data-driven analysis domains.

Participants: Lynn Sutherland, Evie Einstein, Dave Dal Molin, Kevin Wipond, Ken Gamble

A Knowledge-based System for Curriculum Development

Computer Based Training Systems Ltd., a Calgary-based firm, and ARC have undertaken a two-year joint venture in the area of AI and Education. The purpose of the project is to create a knowledge-based system for computer-managed learning. The project is divided into four phases. Phase 1 concerns the development and representation of curricula; in order to assist the user, the system encapsulates extensive expertise in curriculum development and instructional design [8]. Phase 2

incorporates methods for sequencing curricula, developing course maps, and creating and scheduling exams. Phase 3 further extends the work to include the development and maintenance of student models, and comprehensive reporting techniques. Finally, phase 4 implements the statistical analysis of developed curricula. Phase 1 of the system has been prototyped in ART on a Symbolics 3670. Full implementation of the system is underway in C on a Micro Vax II. Alpha testing of the phase 1 prototype has received positive feedback from instructors and curriculum developers. The entire project has benefited from extensive participation by Calgary Board of Education, NAIT, SAIT, Mt. Royal Community College, Red River College, and the Institute of Computer Assisted Learning at the University of Calgary. They have provided much of the curriculum development and instructional design expertise.

Participants: Marlene Jones, Kevin Wipond, Chris Stang

SWIFT: An Expert System for Severe Weather Forecasting

The Alberta Research Council, in collaboration with MacDonald Dettwiler & Associates Ltd., and the Atmospheric Environment Service of Environment Canada is developing a prototype Severe Weather Intelligent Forecasting Terminal (SWIFT). The system will assist in short-term severe weather forecasting and the accurate issuance of severe weather warnings. SWIFT permits integration of conventional severe weather forecasting systems with a new weather representation system. METEOR, one of the expert system modules within SWIFT, relies on severe weather expertise developed at ARC. This rule-based expert system accepts information about meteorology and geography, and produces an assessment of the situation and some advice. The system's knowledge base, consisting of synoptic weather features such as weather fronts and pressure systems, rests on PREFORMA, developed by MacDonald Dettwiler & Associates Ltd. SWIFT demonstrates the integration of large, frequently updated databases with the reasoning capabilities of an expert system. It is implemented on a Xerox workstation, in Interlisp, LOOPS, and OPS4.

Participants: Doug Konkin, Ken Gamble, Julia Driver

In addition to the aforementioned Joint Research Ventures, the Advanced Technologies Department is involved in a number of JRV projects related to Information Technologies, but not directly to artificial intelligence.

4. External Research Associations

The Alberta Research Council encourages researchers in Information Technologies to participate in externally funded research projects outside the Alberta Research Council. In particular, Dr. Marlene Jones is actively involved in three external research projects. She is a member of the Aries Lab at the University of Saskatchewan (see *Canadian Artificial Intelligence*, October, 1987, page 49), where her NSERC grant is used to support joint research with other members of the Aries Lab, including graduate and summer students. In addition, Marlene is a principal investigator on a SSHRCC-funded project concerning the development of an expert system for educational diagnosis. This is joint research with Dr. John McLeod of the Faculty of

Education, University of Saskatchewan. Finally, Marlene, formerly an associate professor at the University of Waterloo (UW), is performing joint research with Dr. Robin Cohen of the University of Waterloo, concerning tailoring explanations to individual users. This work is part of UW's Logic Programming and Artificial Intelligence Group's Theorist project, and is funded by an NSERC strategic grant. Marlene is currently an adjunct professor at the University of Saskatchewan.

Dr. Ernie Chang's external research activities centre around distributed artificial intelligence. Ernie's NSERC grant is used to support graduate and summer students working in the area of Participant Systems. In addition, Ernie is currently an adjunct professor at the University of Calgary.

Additional Information

For information on these or any of the other activities undertaken by Information Technologies please contact:

Betty Ann Snyder
Advanced Technologies Department
Alberta Research Council
3rd floor
6815 8th Street N.E.
Calgary, Alberta
T2E 7H7
Phone: (403)297-2672
e-mail: betty@noah.arc.cdn

References

- [1] Chang, E.J.G., Gamble, K., Masrani, R. and McIlraith, S., "The Process Planner II: An Application of AI to Manufacturing" *Proceedings of IEEE Montech '87 Conference, Compint '87*, October, 1987.
- [2] Chang, E., "Participant Systems" *Distributed Artificial Intelligence*, Huhns, M. (editor). Morgan Kaufmann Publishers. Los Altos, California, 1987. pages 311-339.
- [3] Masrani, R. and McIlraith, S., "ISIM: Towards an Integration of Artificial Intelligence and Simulation" (extended abstract). *Second AAAI Workshop on AI and Simulation* Seattle, Washington, July, 1987.
- [4] McIlraith, S., "Qualitative Modeling: Application of a Mechanism for Interpreting Graphical Data" *Proceedings of CSCSI'88* (to appear), Edmonton, Alberta. June, 1988.
- [5] Chang, E., Clark, D., and Sidebottom, G., "Hermes: Heuristic Emergency Response Management Expert System" *Proceedings of 5th Technical Seminar on Chemical Spills* Montreal, Quebec. February, 1988.
- [6] Gullichsen, E. and Chang, E., "Pattern Classification by Neural Network: An Experimental System for Icon Recognition" *Proceedings of IEEE Conference on Neural Networks* San Diego, California. August, 1987.
- [7] Sutherland, S. and Gamble, K., "A Model for Generating Natural Functions" *Proceedings of Coupling Symbolic and Numeric Computing in Knowledge Based Systems Workshop* Seattle, Washington. July, 1987.
- [8] Wipond, K. and Jones, M., "Curriculum and Knowledge Representation in an Expert System for Curriculum Development" *Proceedings of 1988 International Conference on Intelligent Tutoring Systems* (to appear). Montreal, Quebec. June, 1988.

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	Qualitative Reasoning
Commonsense Reasoning	Temporal Reasoning
Deductive Reasoning	Planning
Diagnostic and	Knowledge Representation Formalisms
Abductive Reasoning	Theories of the Commonsense World
Evidential Reasoning	Theories of Knowledge and Belief
Inductive Reasoning	Belief Management and Revision
Nonmonotonic Reasoning	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)

ADDISON-WESLEY TRAINING SYSTEMS

"We Train the Leaders"

KNOWLEDGE-BASED SYSTEMS IN MANUFACTURING

Carnegie Group, Inc.

AVAILABLE
JUNE 1988

- A 7-hour workshop designed for technical professionals and managers to evaluate the feasibility of developing a knowledge-based system in the manufacturing environment, and to detail the process of system development.
- Complete procedures outlined through the use of case studies and interviews and observations of professionals who have incorporated KBS into their organizations.
- Case studies drawn from Ellwood City Forge, Ford Motor Co., and Digital Equipment Corporation.
- Package includes 60-minute Video Tape, Instructional Guide, and six Workbooks.

KNOWLEDGE REPRESENTATION


Douglas B. Lenat,
Microelectronics and Computer Technology Corporation

AVAILABLE
JULY 1988

- A 15-hour course describing the different methods for representing knowledge in artificial intelligence programs.
- Details how to determine which representation is appropriate for a particular application.
- Twelve AI systems for eight organizations are featured throughout the course: Arthur Andersen, Alcoa Laboratories, BBN Laboratories, ESL (a division of TRW), Foxboro Company, MCC, Monsanto Co. and SRI International.
- Package includes 2 Video Tapes, Instructional Guide and 10 workbooks.



For further information, contact:
Debbi Mark or Cora Thompson, General Books Division
Addison-Wesley Publishers Limited
P.O. Box 580, Don Mills, Ontario M3C 2T8
Telephone (416) 447-5101



Compint '87 at Montech '87: The Computer Aided Technologies

by Suhayya Abu-Hakima

Compint '87 à Montech '87: Les technologies assistées par ordinateur

RÉSUMÉ: De façon générale, la conférence s'est avérée profitable et ce, malgré l'annulation des cours spécialisés et une mauvaise planification des présentations. Montech '87 a accueilli environ 600 personnes. A peu près 200 d'entre-elles ont assisté aux sessions portant sur les systèmes experts et l'IA à Compint '87.

Les cours spécialisés ont été annulés à cause du faible nombre d'inscriptions. De plus, l'ouverture a été retardée de deux heures à cause de l'arrivée tardive de la Ministre des Communications, l'honorable Flora Macdonald. Les présentations portant sur l'IA et sur les systèmes experts n'ont pas été bien planifiées. Les orateurs n'avaient que dix minutes pour présenter leur travaux et pour répondre aux questions. Suite à l'expérience de l'an dernier, les organisateurs avaient prévu que des conférenciers ne se présenteraient pas et ils avaient par conséquent placé un grand nombre de présentations dans chaque session.

Le choix des conférenciers invités a beaucoup contribué à la qualité de la conférence. Parmi les orateurs canadiens, notons les représentants de l'Hydro-Québec, Renato De Mori de McGill, Dennis Dorval de International AI, Murdock McKinnon de CAE et Claude Lajeunesse du CRIM (Centre Recherche Informatique de Montréal). Les orateurs américains étaient Marvin Minsky et Stephen Benton, tous deux du MIT Media Lab. Un message rassurant ressort de la conférence, soit que les universitaires et les chercheurs sont en train de coordonner leurs efforts en IA pour mieux servir les intérêts de l'industrie.

10-13 November 1987, Montreal, Quebec

In evaluating the conference overall one has to rank it as useful despite the cancellation of all tutorials and some poor planning of submitted paper sessions. An estimated 600 delegates attended Montech '87 with about 200 participants in Expert Systems and AI sessions at Compint '87.

All tutorials were cancelled due to low enrollment. The opening was delayed two hours due to the late arrival of the Minister of Communications, the honourable Flora Macdonald. The submitted paper presentation sessions in the AI and Expert Systems tracks were badly planned. Speakers had ten minutes to present as well as to answer questions. Due to past experience the organizers planned for speakers not showing up and placed a large number of speakers in each session — all speakers did show, demonstrating the unfortunate results of overbooking.

An excellent selection of invited speakers contributed significantly to the quality of the conference. The Canadian speakers included several from Hydro-Québec, Renato De Mori from McGill, Dennis Dorval from International AI, Murdock McKinnon from CAE and Claude Lajeunesse from CRIM (Centre Recherche Informatique de Montréal). The American speakers included Marvin Minsky and Stephen Benton, both from the MIT Media Lab. A heartening message from the conference is that Canadian Academics and Researchers are coordinating their efforts to better serve the interests of industry with respect to the emerging techniques in AI.

Suhayya Abu-Hakima is a researcher at the National Research Council in Ottawa.

Expert Systems Sessions

There were a total of five expert system sessions, two invited speaker sessions and three submitted paper sessions. A summary of the key points in each of the five is given below under each one's respective heading.

Invited Speakers

Marvin Minsky's Society of Mind

Minsky spoke for an entertaining hour about many disjoint problems that we must face before unlocking the mysteries of the mind. He opened up by wondering why we have to have expert systems. His response to this was that we need them perhaps so that we may boast or complain about how we have to build them. He then jumped to the topic of why we do not yet have a machine that is as versatile as a person. The reason, he told us, is that obviously we do not yet know how to build machines that have common sense. The machines in AI so far mimic human development in reverse. We first constructed experts to play chess like adults do and we later built a system to analyze children's stories. We have finally recognized that an expert is a generalist with specialized knowledge.

The next problem to contend with is that of knowledge and its representation. Knowledge can be represented by one or several of eight possibilities. There is logic, which according to Minsky is a waste of time since it cannot handle exceptions. Rules are another option which provide us with a method to represent 95% of all human problem solving. Minsky praised Conceptual Dependencies as defined by Schank as simple and useful and lamented that they have unfortunately been ignored by people in AI. He then spoke of natural language

schemes, frames, scripts, semantic nets and hypothesis classification. He stressed that inventing knowledge representation is a field where basic research is lacking. According to Minsky, expert system developers are burying their knowledge representation schemes in proprietary systems since publishing the representation is not as profitable as selling it.

Minsky then discussed negative or anti-knowledge. He spoke of 'bugs' which are essentially exceptions that the system must learn to handle. He also spoke of humour which takes the form of a joke or a stupid act and is stored in a person's negative memory. Negative memory holds knowledge of things one should not do since they are "absurd".

Minsky then arrived at the title of the talk, "The Society of Mind". He offered that the mind is a tremendous kludge. According to him, we have more than one hundred regions that control activity. Each of these is a specialized machine with a unique architecture which can communicate with the other specialized machines, something we are not able to do with today's computers. He also spoke of consciousness as short term memory and that we are mostly unconscious and that it is a transcendental mystery that we have so little consciousness.

Minsky received a range of interesting questions. In reply to my question on handling common sense, he said that in a typical expert system we may have two knowledge representations, each with one thousand associated facts. For common sense however, we need hundreds of cooperating knowledge representations with one thousand facts per representation. He was asked to define understanding, which he flatly refused to do, since he did not see the need for defining the obvious. Minsky was also asked to comment on the fifth generation project. He responded that the Japanese have done well in establishing an AI Lab which is proving to be a good training ground. He also said that they have failed with their main objective to build a smart operating system because they used Prolog. According to Minsky, the use of Prolog is a mistake because it is a language that does not handle exceptions. I took this up with him at the end of his talk and gave him several examples of Prolog's exception handling capabilities. He conceded that perhaps not all Prologs are based on pure logic and that some Prologs can handle exceptions.

Expert Systems at Hydro-Québec

Hydro-Québec had three invited speakers at the expert systems session. The first was Alain Brosseau, the Vice-President of research at their research facility, IREQ. He briefly outlined their effort at IREQ where they have about twenty expert systems under development. The systems are applied to problems of fault diagnosis and advice on maintenance of various equipment. They are using a spectrum of tools ranging from ART to Fortran.

The second speaker was Radu Manoliu, advisor to the Vice-President, who spoke of the strategy that they have adopted as a result of their task force on expert systems. He outlined three strategies that were considered. The first was to be the leader, the second was to follow the leader and the third was to wait and see. The first strategy has been left as a challenge to the AI research community. The third was rejected as the perfect strategy to "miss the boat". The second strategy was adopted and its objective is to harness Hydro-Québec's internal applications expertise and to gain experience

building strategic, lucrative expert systems.

The third speaker was Miguel Marin, a researcher who described the problem of power apparatus diagnosis. He told us that they currently spend \$88 million on this problem. Hence it is a problem of strategic and lucrative importance to them. They are currently aiming to build a set of interworking expert systems to diagnose the various pieces of equipment that fail and advise the operating personnel on coping with the failures.

Submitted Papers

Development and Applications

Six talks were crammed into this one-hour session. Four of them described applications of conventional rule-based expert systems. The first application, related to network maintenance, diagnosed problems between a data network switch and a computer. The second speaker described a system being implemented to advise a user on modeling databases. The third speaker described an advisor that formulates a plan for nursing patient care, based on a spreadsheet view of the patient's state. The fourth speaker gave a very detailed description of an M.I.-based program that checks level sensitive scan design rules for VLSI circuits.

I spoke about RATIONALE, an expert system tool with built-in mechanisms that permit the system to reason by explaining and can be applied to any frame-based diagnostic problem. The sixth speaker addressed expert system maintenance and outlined a tool that associates a checklist with each system hypothesis to insure that modifications in the knowledge base do not backfire.

Computer-Aided Engineering Techniques

The only speaker in this session described the problem of optimization. He spoke of the mathematical view of methodologies, the Engineering view of quick solutions and the need for computers to resolve design and optimization problems.

AI Sessions

Invited Speakers

AI and Diagnosis: Renato De Mori

This talk may have been somewhat disappointing to those who were expecting a general discussion of the problems of diagnosis in AI but was informative to those interested in diagnosis and VLSI.

Knowledge-Systems Applications In Canada: Dennis Dorval, International AI

Dennis Dorval gave a very interesting talk about the state of AI in Canada. He first spoke of the components and players in knowledge-based applications. These ranged from the hardware and the software, to the users, developers and funders of expert systems. He spoke of the shift in expenditure in AI from basic research funding to product development funding. Dorval mentioned the funding organizations for applied research and development, including the National Research Council and PRECARN, established for pre-competitive research and which currently has twenty-two Canadian member companies. He described government supported initiatives such as the matching funds program, the

Ontario Centres of Excellence, the Canadian Institute for Advanced Research (CIAR), the Alberta Research Council and the Canadian Workplace Automated Research Centre (CWARC).

Dorval then described a host of Canadian applications. These included some in engineering, environment, transportation, finance, communications and health care. He left us with the message that Canada has world class knowledge that is waiting to be applied in many areas.

R & D in AI at CAE: Murdock McKinnon

Murdock McKinnon who is the director of research and development at CAE gave an enjoyable speech about the state of AI in his organization. He first described his view of the evolution of AI at CAE. In the 70's, the pre-historical age of AI at CAE, when basic research in AI was underway, CAE was entrenched in Assembler and Fortran. In the mid-70's, the age of ancient history, CAE was investing in personal computers and LANs. In the renaissance period of the early 80's, AI spawned powerful workstations and shells and CAE was busy with ADA. In the modern age of the mid-80's, CAE has finally encountered AI and is actively pursuing exploratory applications.

These applications include flight simulators, flight instructor stations, the 'intelligent helmet' for pilots, a supervisory control system, a compensation and detection system for flight and many space station applications in robotics and manipulators leading into the mid-1990's. His message was that there is an air of excitement at CAE associated with the advent of intelligent systems.

CRIM University-Industry Research Initiative: Claude Lajeunesse

CRIM is a centre established by the Quebec government to contribute to technology in the province and as a result is expected to deliver highly qualified people (at the PhD level) to the workforce. They are expecting to achieve this by identifying research areas and collaborating with member Universities and the private sector to carry out the research. Its member Universities include McGill, Concordia, Université de Montréal and Université du Québec à Montréal. They aim to form a pool of researchers and only work on projects that have a participating Company. Any Company in Canada is accepted as a partner.

Submitted Papers

AI Applications

Six speakers were again crammed into a one hour session. The first speaker described a spelling detector and corrector that was implemented on the Macintosh using Le-Lisp. It classified errors phonetically and grammatically and then used various rules of substitution to correct the errors. The next speaker described a portable natural language system that represents knowledge using semantic nets. It was implemented on an IBM-XT personal computer in TLC-Lisp and the speaker stated his disappointment with his unexpected findings that the lexicon and grammar grew as the system was ported from one application to the next.

The third speaker presented an architecture for a distributed expert system. The speaker following described a rule-based, MYCIN-like approach to verify

VLSI circuits. The fifth speaker gave an approach for tackling the problem of intelligent databases. He suggested the use of natural language front-ends and the addition of rules to current entity-relation descriptions to provide specialized handling. The last speaker had a grand total of four minutes to speak (the other speakers overran their meagre allotment of ten minutes). She simply presented her introductory slide about a diagnostic methodology of constructing databases and invited us all to read the proceedings. This was perhaps the best suggestion that any speaker could have made. Many speakers had prepared thirty minute presentations and were forced to shorten them into ten — some did this by speaking at three times the speed.

Panels: University & Industry Partners in Technology Transfer

One idea that the conference planners should be complimented on was that of having a panel at the conclusion of each day. The three panels were: McGill and IBM Bromont, Concordia and CAE, and Université du Québec à Montréal and Tandem. Each provided an example of collaboration between university and industry and raised issues about the collaboration. The advantages included the funding and the equipment gained by the university from industry as well as first-hand knowledge of real world industrial problems. The advantages to industry included the profitability of applying strategic research and acquiring graduates trained in their particular industry. The disadvantages included the question of industry's tight schedules and requirements for confidentiality of research results. Overall, the panels served to highlight the emerging collaboration between industry and university for the advancement of applied research.

Summary

In summary, the planners of the submitted paper sessions should re-examine their strategy of dealing with no-shows. Inconveniencing people who have come across continents prepared to speak for half an hour by cutting them off after seven minutes is deplorable. I suggest they halve the number of accepted papers or they increase the number of parallel sessions.

The organizers of the invited speakers are to be congratulated for their selection. Of particular note was the session on Computer Synthesized Holography by Stephen Benton from the MIT Media Lab. Overall, the conference gave one a good feeling for AI in Canada with particular emphasis on Quebec. The submitted papers were mostly from universities across Canada, with a couple of papers from universities in France and Germany. There was a surprisingly low participation by Canadian industry. However, the panels showed some notable cooperation between researchers, academics and industrialists. This cooperation may stem from industry feeling that they need to catch up. The basic research presented by the universities consisted mostly of the re-implementation of MYCIN for various applications. However, other researchers in Canada are currently addressing important areas in knowledge-based systems such as knowledge representation, reasoning and expert system maintainability and hopefully they will contribute to the next conference.

The Second Conference on Applied Natural Language Processing

by Jugal Kalita

Deuxième conférence au sujet du traitement pratique de langues naturelles

RÉSUMÉ: Le thème des articles présentés à cette conférence était l'application de techniques TLN à des problèmes concrets. La conférence comprenait des ateliers, des présentations d'articles, des étalages commerciaux et des démonstrations d'applications touchant au TLN dont les interfaces homme-machine, les entrées vocales, la production de textes, la traduction automatique, la bureautique, les aides aux écrivains et les didacticiels informatisés.

9-12 February 1988, Austin, Texas

The second conference on Applied Natural Language Processing, held under the sponsorship of the Association for Computational Linguistics, was attended by more than three hundred people from all over North America as well as Europe and Asia. More than a hundred papers were submitted by researchers from over 12 countries and less than one-third of the papers were accepted for presentation at the conference.

The theme of the papers presented was the application of NLP techniques to real world problems. There were tutorials, panel discussions, company exhibits and demonstrations of applications. Papers were presented in a wide range of areas such as human-machine interfaces (including databases, expert systems, and report writers), speech input and output, information retrieval, text generation, machine translation, office automation, writing aids, computer-aided instruction, and tools for NLP. Emphasis was placed on completed rather than intended work.

One notable fact about this conference was that about 60% of the papers presented were from various industries and research organizations. The rest of the papers were by University researchers and students.

Several tutorials were given by noted contributors to their respective fields. These included tutorials on introductory computational linguistics (Allen), machine-readable lexicons (Bogwarev and Levin), spoken language systems (Roncos), NL interface techniques (Hafner), Logic for NLP (Moore), and machine translation (Nirenburg).

There were several papers discussing aspects of natural language generation. Rosner discussed issues encountered in the design and implementation of the generation module of the German-Japanese translation system, SEMSYN. McDonald and Meteor presented a method for building a NL interface from underlying application programs to the linguistic realization component, MUMBLE-86, taking into consideration the linguistic naivety of most application programs. Boggess' discussion of algorithms to facilitate text production for handicapped individuals showed that NLP may one day be beneficial to disadvantaged sections of the society.

Papers on natural language systems included Brennan's paper on a domain independent strategy for multi-media articulation of answers in a NL database interface, and Hayes' discussion of a commercially

available rule-based system for categorization of news stories developed at the Carnegie Group at Pittsburgh. Several papers discussed various issues involved in the development of the LUCY system released to shareholders by MCC of Austin. Rich and LuperFoy's paper dealt with an integrated strategy for anaphora resolution, whereas Wittenburg and Burnett critically evaluated the "conduit" model for NLP as used in LUCY.

Handling of syntactic and semantic ambiguities were also discussed in several papers. Hurum's paper presented domain-independent heuristics for resolution of "scope ambiguities", Bear and Hobbs dealt with localization of resolution methods for syntactic ambiguities, and Newman adapted heuristics-guided A* search for resolution of various types of ambiguities.

Handling of ill-formed inputs was also discussed at length at this conference. Van Berkel and De Smedt presented details of a system that corrects word-level errors by using a novel triphone analysis, and Means discussed a similar heuristic-based system for recognition of word-level errors. Richardson and Braden-Harder described their experience in developing the CRITIQUE system at IBM which identifies word-level as well as stylistic errors in documents, and suggests appropriate changes.

Machine-readable dictionaries generated substantial amount of interest in this conference. Among the notable papers was by Neff et. al discussing a system for grammatically describing and parsing entries from machine-readable dictionary tapes. Another was by Jensen and Binot describing techniques for utilizing information in the text of such lexicons to aid various aspects of natural language processing such as attachment of prepositional clauses and resolution of pronoun reference. Fox et. al presented a paper discussing details of building an on-line thesaurus in the form of a large semantic network.

The paper on the current status of the machine translation system EUROTRA (translating among nine official languages of the EEC) by Varile and Lau threw light on the various issues confronted by the designers and the implementors. Gebruers described a valency model developed within the Belgian METAL MT project for enhancing modularity and multilinguality.

Among other papers which came to my attention was that of Church where he discussed a program which uses principles of linear time dynamic optimization programming to find assignments of parts of speech to words. Another interesting paper was Ejherd's comparison of a regular expression method and a stochastic method of parsing natural language with

Jugal Kalita is a PhD student at the University of Pennsylvania; Jugal holds a masters degree from the University of Saskatchewan.

respect to their successes in identifying basic clauses in unrestricted English text. Allport's paper on summarization of free-form text by identifying interesting events, and Kalita and Shende's system for modeling a set of processes encountered in a typical office environment, and planning and generating impressive paragraph-level text were also very interesting. Slocum's discussion of the morphological processor for the NABU system which can handle eight languages (Arabic, Chinese, English, French, German, Japanese, Spanish and Russian) was stimulating in spite of being general in its contents. Rau and Jacob's paper on gracefully integrating various sources of knowledge for text processing was also well received.

In addition to the paper presentations, there was a lively panel discussion where several leading innovators in commercial natural language interfaces participated. Among the participants were AI Corporation's Harris (developer of Intellect), Symantec's Hendrix (Q&A), Bobrow of BBN (Parlance), Ginsparg of Natural Language Inc. (DataTalker) and Klein of Singular Solutions Engineering (Lotus HAL). They identified

where their systems have succeeded, where they are less than successful and where they are heading. They emphasized the need for integrating natural language into other means of interacting: pointing, drawing and menus. Additionally, the consensus of the panel was that in order to succeed, natural language products must achieve all traditional data processing qualities, including efficient processing, good documentation, trustworthy support, and compatibility between releases.

There were several exhibits and demonstrations. Among the interesting demonstrations were AT&T's TELI natural language question answering system, CMU's universal parser/compiler for machine translation, Logicon's DBG message understanding system, MCC's LUKE knowledge editor, SRI's TACITUS discourse understanding program, SEMSYN (West Germany) group's GEOTEX text generator, UNISYS's PUNDIT text processing system, IBM's CRITIQUE system, TNO Institute's (the Netherlands) information retrieval and spelling correction system, and Planning Research Corporation's PAKTUS NLP development environment system.

University of Toronto Department of Computer Science

Lisp programming positions available at the University of Toronto

Two programming positions are available within the Artificial Intelligence group of the University of Toronto. The positions are funded for a three year period for a research project aiming at the development of an expert system shell. Candidates for the positions are expected to have an M.Sc. or Ph.D. degree in computer science and to have good knowledge of Lisp and the Lisp machine computing environment. Salary range for the positions is \$32,000-\$53,000 depending on qualifications. For more details contact John Mylopoulos at

Department of Computer Science
University of Toronto
10 King's College Road
Toronto, Canada M5S 1A4

tel: (416) 978-5180
email: jm@ai.toronto.edu

University of Waterloo Department of Computer Science

The Department of Computer Science at the University of Waterloo comprises 40 full-time faculty members engaged in research and teaching activities. The government of the Province of Ontario has awarded a five-year (renewable) centre of excellence in information technology to the University of Waterloo and the University of Toronto (with participation from Queen's University and the University of Western Ontario). The Department of Computer Science and the Institute for Computer Research are housed in the William G. Davis Computer Research Centre, a new building with 300,000 sq. ft. dedicated to computer research laboratories.

The University of Waterloo invites applications for faculty positions in Computer Science. A Ph.D. in computer science is required, with evidence of outstanding research accomplishment or potential. All areas will be considered; candidates with research interests in artificial intelligence, hardware and software systems, or symbolic computation are strongly encouraged. Salary is commensurate with experience. Applications from women candidates and recent Ph.D. graduates are particularly welcome. Inquiries should include a curriculum vitae and the names of three references and should be directed to the chairman: Prof. John A. Brzozowski, Department of Computer Science, University of Waterloo, Waterloo, Ontario, Canada N2L 3G1, brzozo@water.waterloo.edu.

In accordance with Canadian Immigration requirements, this advertisement is directed to Canadian Citizens and Permanent Residents.

VRP* EMBEDDED REAL-TIME EXPERT SYSTEMS

V is for Versatile:

The Autometrics* VRP Expert System Developer's Kit is for implementation of fast, compact expert systems embedded in application software, system software, or firmware for almost any computing environment - microcontrollers to mainframes.

- Backward and forward chaining
- Propositional inference
- Uncertainty management

R is for Rule-based:

Build and test a knowledge base of plain-language IF-THEN rules on a PC, then transfer it to the target system. The VRP Expert System Developer's Kit contains:

- VRP Expert System Developer's Guide
- Application Notes, examples
- Rulebase compiler and debugger
- Optional target runtime modules

P is for Portable:

Link the VRP inference procedure and runtime support routines into your application software and install with the rulebase in the target system.

- Re-entrant, ROMable, and fast
- Easy software and I/O interfacing
- Z80, 8085, 8086, 68000, and others

VRP Applications:

Data interpretation; self-configuring systems; diagnostic or advisory functions; process control; natural language processing; performance monitoring; intelligent device drivers.

Autometrics is for industrial AI.

AAL Autometrics Associates Ltd.
P. O. Box 6214 - Station C
Victoria, B. C., Canada V8P 5L5
Telephone (604) 652-2679

*Autometrics and VRP are trademarks of AAL Autometrics Associates Ltd.

Book Reviews

edited by Graeme Hirst

Knowledge Systems and Prolog: A Logical Approach to Expert Systems and Natural Language Processing

Walker, Adrian (editor); McCord, Michael ; Sowa, John F.; and Wilson, Walter G.

[IBM T.J. Watson Research Center and IBM Systems Research Institute]

Reading, MA: Addison-Wesley, corrected reprinting, March 1987, xii+475 pp.

Hardbound, ISBN 0-201-09044-9, Cdn\$43.35

*Reviewed by
Innes Ferguson
Bell-Northern Research*

The aim of this book is to provide readers with an introduction to Prolog, from both a theoretical and a programming point of view. In particular, two areas of Prolog use are stressed: expert systems and natural language processing programs, or more simply, knowledge systems. To this effect, the book describes the interplay between both principles and practice of knowledge systems — it is the belief of the authors that Prolog, with its formal foundations in logic, can bridge the gap between the two approaches to studying knowledge systems. In a sense, the book calls for empiricists and formalists to consider Prolog and several of its related techniques and systems.

Depending on the reader's own area of interest, the book can be used in several different ways. Clear instructions are given as to which sections should be read if your interests lie in (i) the possible uses of expert systems, (ii) natural language processing, (iii) programming in Prolog, or (iv) the logical fundamentals of the authors' approach to knowledge systems and Prolog. It is fair to say that such areas of interest are addressed in the book, but I'm not sure how much the reader would get out of the book by only reading the recommended sections. It would seem to me that the "big picture" (i.e. the relationship between formal and empirical approaches to knowledge system modeling) would be lost.

Of course, losing the "big picture" may not be as bad as it sounds. In fact, the authors themselves suggest that the book could simply be used either as a textbook to learn Prolog for the first time, or to add advanced programming techniques to a reader's prior knowledge of Prolog. Personally, however, I would not recommend this book as a primer for programming in Prolog. While several interesting techniques are demonstrated throughout the book (e.g., meta-interpretation, graph searching techniques, and inferencing methods), the tutorial section on Prolog is not thorough enough — familiarity with many basic concepts is assumed and no clear order appears to exist for introducing new concepts.

There are four major parts to this book. The first of these provides the reader with an introduction to the Prolog language itself. This is performed in a problem-oriented fashion, providing the user with many useful example programs, as well as several unsolved exercises. Much of what is written in this part of the book (Chapters 1 and 2 and Appendix A) was derived from course material used by the authors at IBM's various teaching centres. The particular syntax used corresponds to IBM's own Prolog, so this book can also be used as a supplement to the IBM Prolog programming manual. A clear distinction between pure or declarative Prolog and procedural Prolog is provided, along with the pros and cons of each style. Numerous "good programming" techniques are described, along with several interesting programs for implementing meta-interpreters, debuggers, proof trees, and graph searching techniques.

The second part (Chapter 4) deals with expert systems design in Prolog. Several different knowledge representation techniques are described, including rules, frames, and "pure" logic. Similarly, several inference mechanisms are described and exemplified, including backward and forward chaining and inheritance. One of IBM's own expert system shells, Syllog, is described together with an example knowledge base for manufacturing planning. While the description of the shell itself is interesting, I feel a better example could have been chosen to demonstrate Syllog's more interesting capabilities, which include assisting the programmer in the acquisition of domain knowledge, and reasoning directly from base facts and syllogisms.

The third part (Chapter 5) addresses natural language processing. In particular, two formalisms, McCord's Modular Logic Grammars (MLGs) and Definite Clause Grammars (DCGs), are introduced and compared, together with several mechanisms for representing lexicons, writing rule compilers for MLGs and DCGs, and techniques for handling certain "interesting" syntactic constructs in English. This chapter also considers natural language database querying, and toward the end, comes full circle by investigating various procedures which produce logical form (semantic) analyses from given syntactic analyses.

The last part (Appendix B plus several sections throughout the text) is not very extensive, but is certainly quite interesting. Here a logical basis for Prolog using model theory is given, with special attention paid to the meaning of negation in Prolog, and how it relates to this theory. This theory is then applied to the problem of writing efficient — as well as semantically sound — language interpreters.

In summary, this book, which is about AI, logic, and programming, attempts to bridge the gap between knowledge systems and the formal foundations of reasoning in logic. I strongly recommend it to those readers who are interested in advanced Prolog programming techniques for implementing expert systems and natural language processing programs. Many examples (both solved and unsolved) are given, including many excerpts of code that would be extremely useful for building or studying one's own knowledge systems. I'm not convinced, however, that this would be a good introduction to programming in Prolog. I would

recommend that the reader have some prior knowledge of the language before reading this book. As far as attempting to bridge the gap between the empirical and formal views of knowledge systems, I believe that the authors' approach is correct, and that this text is definitely a good step in the right direction.

Innes Ferguson is a Research Scientist with the AI Exploratory Group at Bell-Northern Research, Ottawa.

Expert systems: Tools, Techniques, and Applications

Klahr, Philip; and Waterman, Donald A. (editors)
[The Rand Corporation]

Reading, MA: Addison-Wesley, 1986, vi+441 pp
Hardbound, ISBN 0-201-14186-8, Cdn\$35.95

Reviewed by
Peter Davies
Expert Solutions

This volume is a collection of papers on recent work at the Rand Corporation in artificial intelligence research and development. An introductory chapter details the history of AI at Rand. The remainder of the book is split into three sections, entitled (not surprisingly) Tools, Applications, and Techniques.

Almost without exception, the articles in the book are interesting and well written. However, I had some difficulty determining its intended audience. The neophyte, having read (the late) Donald Waterman's previous publication, *A guide to expert systems*, may assume this book to be the next step in understanding the field. In fact, it is several steps down the road. The more knowledgeable reader, having some exposure to expert system development, may be frustrated to read about tools to which they have no access and a limited array of application areas. However, those interested in the application of AI to military systems should find plenty of useful information.

The introduction reads like a who's-who of artificial intelligence, computer science, and mathematics. One gets the impression all the leading lights have been sometime employees or consultants at Rand. To name only a few: Von Neumann, Newell, Simon, Shaw, Feigenbaum, Tarski, and Kleene. (And you thought they were all just academics!)

The Tools section discusses ROSIE and ROSS, both of which were developed, and are used, at Rand. ROSIE is an expert system development tool which combines a rule-based representation with a procedure-oriented language design. The narrative covers all aspects of the tool, including its history, and makes interesting reading. It offers insights into the kinds of features a good environment requires to aid development of large-scale, complex systems, providing useful information for builders of tools and applications. I was a little bemused, however, to find in an article at this level, an introduction to expert systems. Surely the authors were not aiming this article at people without any background in the field! ROSS is an object-oriented language for constructing simulations. Although much shorter than the chapter on ROSIE, this chapter gives a good flavour of the language. It also provides an introduction to object-oriented programming and how these techniques can be applied to

simulation. After reading about these two tools, some readers may be disappointed to learn that, as far as I know, neither tool is available as a product.

Since everything discussed in this book was done at Rand, the range of applications covered is, necessarily, limited. Three were funded through DARPA and focus on international terrorism, tactical air targeting and the simulation of tactical warfare. This will be unsettling to many, and is a reflection of one of the major reasons behind the massive interest and funding the AI field has enjoyed in recent years south of the border. The other application discussed is concerned with legal decision-making. All four chapters are very well-written, in-depth discussions of the respective systems. They also mesh well with the tools section. Two were developed in ROSIE, one in RITA, a forerunner of ROSIE, and the remaining application was written using the ROSS language.

The section on techniques is, for my money, the most rewarding. It deals with some important and interesting research issues. The first two chapters are concerned with knowledge acquisition. Chapter 8 discusses the EP-2 system, which builds systems from examples. The techniques and ideas employed are far beyond the inductive "shells" we see in the marketplace today. Chapter 9 describes how programs can be constructed which refine their own knowledge bases, on the basis of their performance. I found this to be the most interesting chapter, providing insights on machine learning. It also contains a nice overview of prior learning systems research. Chapter 10 deals with plausible inference from uncertain knowledge, another area still in the research stage. The system discussed, INFERNO, uses a probabilistic approach and advises the user on ways to adjust for inconsistencies it encounters. Technically speaking, this is probably the most difficult chapter in the book. However, this is to be expected in an area where the imprecision of human decision-making is modeled using the precise language of mathematics. The last chapter concerns distributed problem solving and discusses this in the domain of Air Fleet Control.

I can certainly recommend this book to those interested in military applications of artificial intelligence. Also, anybody currently involved in development of applications in other domains should find plenty of useful information concerning current and research issues. It is not recommended for those with little exposure to the field.

Peter Davies is a consultant, developer and researcher in expert systems and other AI applications. He is the principal of Expert Solutions, a Toronto-based AI company.

**Artificial intelligence and Tutoring Systems:
Computational and Cognitive Approaches to the
Communication of Knowledge**

Wenger, Etienne
[University of California, Irvine]

Los Altos, CA: Morgan Kaufmann, 486 pp
Distributed in Canada by John Wiley
Hardbound, ISBN 0-934613-26-5; US\$29.95

Reviewed by
Jim E. Greer
University of Saskatchewan

Although artificial intelligence and tutoring systems is a

relatively new research discipline, spanning scarcely more than a decade and a half, it has experienced rampant growth. This book attempts, with a large measure of success, to organize all of intelligent tutoring systems (ITS) research within a unifying framework, and to use the framework to evaluate results and contributions of major ITS research efforts. Wenger presents an up-to-date overview of the major projects in the area, providing a comprehensive bibliography covering publications through 1987.

The book is divided into three parts. The first part provides a brief introduction to "knowledge communication," a framework that Wenger introduces as an organizational structure for the book and the ITS research discipline. The basic issues of representing domain and pedagogical knowledge, of modeling student knowledge and behaviour, and of designing interfaces are introduced within the knowledge communication framework.

Part Two is a lengthy (nearly 300-page) account of major ITS research projects. It is organized into eleven chapters around four apparent general themes:

1. the evolution of mental models and qualitative reasoning
2. the development of learning environments
3. the diagnosis of student misconceptions
4. planning and delivery of instruction

Within each theme various systems are discussed with emphasis on their primary research contributions. Although it is impossible to discuss each system in great detail, Wenger has successfully distilled out the important elements of each system, and has provided numerous references and interesting bibliographic notes to guide the interested reader deeper into the field. Many ITS projects have made contributions across the four themes and more generally to theories and techniques fundamental to artificial intelligence and cognitive science research. Wenger's system-by-system account occasionally leads to descriptions of contributions incompatible with the organization of this part of the book. Overall, Part Two represents a noble effort at organizing the diverse approaches to intelligent tutoring, one which no one else has yet presumed to undertake.

The third part of the book consists of a survey of the current state of ITS research along another dimension, that is with respect to typical ITS components. While Part Two might be thought of as a horizontal pass through ITS systems, Part Three is a vertical pass through the issues pertaining to instructional domains, student modeling, diagnosis and didactics. This part of the book elaborates and attempts to justify Wenger's knowledge communication framework for ITS research introduced in Part One.

Parts of the book are tersely written, summarizing in a loaded sentence or two, concepts that may deserve deeper elaboration. Some sections also tend to contain excessive jargon. Nevertheless neither of these concerns is serious enough to detract significantly from the value of the book.

Although there are a number of other new books on intelligent tutoring systems, *Artificial Intelligence and Tutoring Systems* by Etienne Wenger can be highly recommended on the basis of its broad coverage and insightful organization of ITS research. It points toward the development of a research methodology for the design and evaluation of ITS research. This book is a particularly timely contribution since the development of

acceptable methods and standards for ITS research (and AI research in general) deserves serious consideration.

The book serves as an ideal ITS textbook provided that the reader has some background in artificial intelligence and cognitive science. It is much more comprehensive than the edited collections of papers that were previously available.

Jim Greer is a Postdoctoral Fellow in Computational Science at the University of Saskatchewan. He recently completed a PhD from the University of Texas at Austin. His research interests are in the area of intelligent educational systems and cognitive science. He is a co-founder of the Laboratory for Advanced Research in Intelligent Educational Systems (ARIES) at the University of Saskatchewan.

Briefly Noted

Thinking Machines: The Evolution of Artificial Intelligence

Pratt, Vernon
[Lancaster University]

Oxford: Basil Blackwell, 1987, xii+254 pp
Distributed in Canada by Oxford University Press
Hardbound, ISBN 0-631-14953-8, Cdn\$53.95

Pratt is not an AI researcher but a historian and philosopher of science. His purpose in this book is to present the history of the idea of the mechanization of thought, commencing with the rise of modern mathematics from the Middle Ages on. There are three featured players: Leibniz, Babbage, and Turing. Pratt treats each in great detail, with diagrams explaining the workings of the various calculating and computing machines that they devised. Modern, post-Dartmouth AI is only a small part of the book.

The book should not be taken as a clone of those by Haugeland* or McCorduck**. Haugeland is a philosopher, but his excellent book is more interested in explicating the ideas of AI and the underlying concepts of mind than explaining their history. McCorduck's history of the field concentrates on the recent years. There are also many other books on the history of computing, but I know of none with Pratt's emphasis on the ideas of AI.

Pratt's book will be of interest to anyone in AI who wants to know more of the intellectual history of the field. And it should be required reading for those all-too-numerous bores who think AI came out of nowhere about three years ago, that it has no significant history, and that anything anyone does in AI is ipso facto unprecedented.

—G.H.

L'échéance pour le numéro de juillet est le 15 mai

*John Haugeland, *Artificial intelligence: The very idea*. The MIT Press, 1985.

**Pamela McCorduck, *Machines who think*. Freeman, 1979.

The Articulate Computer

McTear, Michael
[University of Belfast]

(The language library)
Oxford: Basil Blackwell, 1987, ix+240pp.
Distributed in Canada by Oxford University Press
Hardbound, ISBN 0-631-14009-3, Cdn\$90.75

Michael McTear's book *The Articulate Computer* will be useful to anyone in or out of AI who wants to know more about natural language processing. Although the book is an introduction to the field, it is not a textbook. Rather, it is addressed to readers in neighbouring disciplines, and, indeed, the intelligent lay readers, who want to know more about natural language processing, its successes, its issues, and its problems, without the technical density or pedagogical style of a text book. That is not to say that McTear shies away from technical detail; on the contrary, the book contains plenty of it. But it is detail meant to produce thorough understanding, not to teach the reader to write their own programs. Only a minimal knowledge of linguistics and computing is assumed.

McTear's presentation is comprehensive, including syntax, semantics, issues in knowledge representation, and problems in extended discourse — especially cooperative user interfaces. This book would be an excellent starting point for anyone in other areas of AI who wants to know more about natural language processing but feels hesitant about tackling the technical literature straight off. Unfortunately, the amazing price, for which there is no apparent reason, will deter would-be purchasers.

—G.H.

A Guide to Commercial Artificial Intelligence

Rauch-Hindin, Wendy B.
[Hi-Tech Editorial and Consulting, Inc]

Englewood Cliffs, NJ: Prentice-Hall Inc, 1988,
xvii+523 pp
Paperback, ISBN 0-13-368770-8, US\$28.95

A revised and slightly abridged version of the author's earlier two-volume set *Artificial intelligence in business, science, and industry*, this book is a comprehensive and detailed applications-oriented introduction to AI, with an emphasis on technical details, case histories, and commercial products.

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.

‡ Applications of Expert Systems: Based on the Proceedings of the Second Australian Conference *Quinlan, J Ross*

[New South Wales Institute of Technology]
Glasgow and Sydney: The Turing Institute Press in association with Addison-Wesley, 1987, xvi+223 pp
Hardbound, ISBN 0-201-17449-9, US\$24.95

Artificial Intelligence Programming [2nd edition]

Charniak, Eugene; Riesbeck, Christopher K.; McDermott, Drew V.; Meehan, James R.
[Brown University; Yale; Yale; Cognitive Systems Inc]
Hillsdale, NJ: Lawrence Erlbaum, 1987, xvi+533 pp
Distributed in Canada by Copp Clark Pitman Ltd.
Hardbound, ISBN 0-89859-609-2, Cdn\$46.95

‡ The Fifth Generation Fallacy: Why Japan is Betting its Future on Artificial Intelligence

Unger, J Marshall
[Dept East Asian Languages and Literatures, University of Hawaii, Honolulu]
Oxford University Press, 1987, x+230 pp
Hardbound, ISBN 0-19-504-939-X, Cdn\$37.50

An Introduction to Splines for Use in Computer Graphics and Geometric Modeling

Bartels, Richard H.; Beatty, John C.; and Barsky, Brian
[University of Waterloo; University of Waterloo; University of California, Berkeley]
Los Altos, CA: Morgan Kaufmann, 1987, xiv+476 pp
Hardbound, ISBN 0-934613-27-3, US\$38.95

Knowledge Aided Information Processing

[Essays dedicated to Prof Freidhart Klix]
van der Meer, Elke; and Hoffmann, Joachim (editors)
[Department of Psychology, Humbolt University and Department of Psychology, Academy of Sciences of the German Democratic Republic]
Amsterdam: North-Holland, 1987, xii+324 pp
Hardbound, ISBN 0-444-70276-8, Dfl 190.00

The Knowledge Frontier: Essays in the Representation of Knowledge

Cercone, Nick; and McCalla, Gordon (editors)
[Simon Fraser University and University of Saskatchewan]
(Symbolic computation series: artificial intelligence)
NY: Springer, 1987, xi+512 pp
Hardbound, ISBN 0-387-96557-2, Cdn\$58.50

The Little LISPer (trade edition)

Friedman, Daniel P.; and Felleisen, Matthias
[Indiana University]
Cambridge, MA: The MIT Press, 1987, xiv+186 pp
Paperback, ISBN 0-262-56038-0, US\$12.95

‡ 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

Productive Prolog Programming

Schnupp, Peter; and Bernhard, Lawrence W.
[Interface GmbH and Soft & Hard Software]
Prentice-Hall, 1987, xiv+296 pp
Paperback, ISBN 0-13-725110-6

Readings in Human-Computer Interaction: A Multidisciplinary Approach

Baecker, Ronald M.; and Buxton, William A. S.
(writers and editors)
[University of Toronto]
Los Altos, CA: Morgan Kaufmann, 1987, xiii+738 pp
Paperback, ISBN 0-934613-24-9, US\$32.95

‡ Readings in Non-Monotonic Reasoning

Ginsberg, Matthew L.
[Stanford University]
Los Altos, CA: Morgan Kaufmann, 1987, 481 pp
Paperback, ISBN 0-934613-45-1, US\$26.95

‡ Reasoning about Change: Time and Causation from the Standpoint of Artificial Intelligence

(The MIT Press series in artificial intelligence)
Shoham, Yoav
[Department of Computer Science, Stanford
University]
Cambridge, MA: The MIT Press, 1988, xvi+200 pp
Hardbound, ISBN 0-262-19269-1, US\$25.00

Special Issue on Artificial Intelligence

Daedalus 117 (1), Winter 1988
(Journal of the American Academy of Arts and
Sciences)

‡ Structured Induction in Expert Systems

Shapiro, Allen D.
[Cogensys Corporation]
Glasgow and Wokingham: The Turing Institute Press in
association with Addison-Wesley, 1987, xiii+134 pp
Hardbound, ISBN 0-201-17813-3

‡ 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

Computational Intelligence

Abstracts for 3(4) November 1987
Special Issue on Machine Learning
Guest Editor: Larry Rendell

A Discussion of a Report by Ehud Shapiro

Ranan B. Banerji
Saint Joseph's University
Philadelphia, Pennsylvania

The paper is an annotated summary of Ehud Shapiro's report "The Induction of Theories from Facts." In the view of the author, Shapiro's work forms a very good foundation for work in the field of learning. It gives a clear definition of the term "learning" in a way which is

both intuitively acceptable and renders learning algorithms amenable to precise analysis. It also establishes a paradigm for learning algorithms which is precise enough that it can serve as a benchmark for future development as well as for the analysis of presently available algorithms.

Learning to Control a Dynamic Physical System

Margaret E. Connell and Paul E. Utgoff
Department of Computer and Information Science
University of Massachusetts
Amherst, Massachusetts

This paper presents an approach to learning to control a dynamic physical system. The approach has been implemented in a program named CART and applied to a simple physical system studied previously by several researchers. Experiments illustrate that a control method is learned in about 16 trials, an improvement over previous learning programs.

A Foundational Approach to Autonomous Knowledge Acquisition

James P. Delgrande
School of Computing Science
Simon Fraser University

A formal foundational approach to autonomous knowledge acquisition is presented, together with its relation with logical knowledge representation systems. For knowledge acquisition, it is assumed that the only information available is a stream of facts, or ground atomic formulae. On the basis of this information, hypotheses expressed in set-theoretic terms and concerning the application domain may be proposed. As further instances are received, the hypothesized relations may be modified or discarded, and new relations formed. Formal systems are derived by means of which the set of potential hypotheses is precisely specified, and a procedure is derived for restoring the consistency of a set of hypotheses after conflicting evidence is encountered. This framework is intended both as a basis for the development of autonomous systems that learn from examples and as a neutral point from which such systems may be viewed and compared.

The relation between acquiring knowledge and reasoning deductively is also explored. The hypothesis formation process is considered as a source of sentences for a first-order system that can reason with conjecture and knowledge, and this deductive component is considered as a source of "a priori" knowledge for the hypothesis formation process. However, the interaction of these approaches leads to immediate difficulties which likely require informal, pragmatic techniques for their resolution.

Learning in Artificial Neural Systems

William E. Hohensee and Christopher J. Matheus
Inductive Learning Group
Department of Computer Science
University of Illinois at Urbana-Champaign

This paper presents an overview and analysis of learning in "Artificial Neural Systems" (ANS's). It begins with a general introduction to neural networks and connectionist approaches to information processing. The

basis for learning in ANS's is then described, and compared with classical machine learning. While similar in some ways, ANS learning deviates from tradition in its dependence on the modification of individual "weights" to bring about changes in a knowledge representation distributed across connections in a network. This unique form of learning is analyzed from two aspects: the selection of an appropriate network architecture for representing the problem, and the choice of a suitable learning rule capable of reproducing the desired function within the given network. The various network architectures are classified, and then identified with explicit restrictions on the types of functions they are capable of representing. The learning rules, i.e. algorithms that specify how the network weights are modified, are similarly taxonomized, and where possible, the limitations inherent to specific classes of rules are outlined.

A Computational Theory of Motor Learning

Wayne Iba and Pat Langley

Department of Information and Computer Science
University of California, Irvine

In this paper we present a computational theory of human motor performance and learning. The theory is implemented as a running AI system called MAGGIE. Given a description of a desired movement as input, the system generates simulated motor behaviour as output. The theory states that skills are encoded as "motor schemas" which specify the positions of a limb at selected points in time. Moreover, there exist two natural representations for such knowledge: "viewer-centred" schemas describe visually perceived behaviour, while "joint-centered" schemas are used to generate behaviour. When the model acts upon these two representational formats, they have quite different behavioural characteristics. MAGGIE performs the desired movement within a feedback control paradigm, monitoring for errors and correcting them when it detects them. Learning involves improving the joint-centred schema over multiple practice trials; this reduces the need for monitoring. The model accounts for a number of well-documented motor phenomena, including the speed-accuracy trade-off and the gradual improvement in performance with practice. It also makes several testable predictions. We close with a discussion of the theory's strengths and weaknesses and directions for future research.

A Reasoning-Based Approach to Machine Learning

Krish Purswani and Larry Rendell

Department of Computer Science
University of Illinois at Urbana-Champaign

This paper describes a novel approach to machine learning, based on the principle of learning by "reasoning". Current learning systems have significant limitations such as "brittleness," i.e. the deterioration of performance on a different domain or problem and lack of power required for handling real-world learning problems. The goal of our research was to develop an approach in which many of these limitations are overcome in a unified, coherent and general framework. Our learning approach is based on principles of

reasoning, such as the discovery of the "underlying principle" and the recognition of the deeper basis of similarity, which is somewhat akin to human learning. In this paper, we argue the importance of these principles, and tie the limitations of current systems to the lack of application of these principles. We then present the technique developed, and illustrate it on a learning problem not directly solvable by previous approaches.

Similarity-Based Learning and its Extensions

Larry Rendell

Department of Computer Science
University of Illinois at Urbana-Champaign

This paper synthesizes a number of approaches to concept representation and learning in a multi-layered model. The paper emphasizes what has been called "similarity-based learning" from examples, although this view is extended to address wider issues. The paper pays particular attention to requirements for incremental and uncertain environments, and to interrelationships among concept purpose, concept representation, and concept learning.

One goal of the paper is to unite some of the notions underlying recent research, in an attempt to construct a more complete and extensible framework. This framework is designed to capture representations and methods such as those based on hypothesis search and bias selection, and to extend the ideas for greater systems capability. This leads to a specific perspective for multi-layered learning which has several advantages, such as greater clarity, more uniform learning, and more powerful induction.

The approach clarifies and unifies various aspects of the problem of concept learning. Some results are: (1) Various concept representations (such as logic, prototypes, and decision trees) are subsumed by a standard form which is well suited to learning, particularly in incremental and uncertain environments. (2) Concept learning may be enhanced by exploiting a particular phenomenon in many spaces — this phenomenon is a certain kind of smoothness or regularity, one instance of which underlies the similarity in SBL systems. (3) The paper treats the phenomenon in a general way and applies it hierarchically. This has various advantages of uniformity. For example, the model allows layered learning algorithms for concept learning all to be instantiations of one basic algorithm. A single kind of representation (an instantiation of the standard form) is prominent at each level. The combination of representation and algorithm allow fast, accurate, concise, and robust concept learning.

Inductive Ambiguity and the Limits of Artificial Intelligence

Satosi Watanabe

University of Hawaii
Honolulu, Hawaii

Most artificial intelligence systems are based fundamentally on deduction. Such systems usually rely on a number of deductive heuristics. We can justify this reliance only when we can assume that the logical lattice of propositions is closed. In particular, this restriction forces us to concede that deductive AI systems cannot be capable of performing activities such as abduction and inductive evaluation, which are at the heart of scientific

activity. To illustrate, we will look at diagnostic practices in medicine and compare the inductive-scientific and deductive-technological processes. Deductive AI systems can help us only with the latter type of analysis.

Abstracts for 4(1) February 1988 "Taking Issue"

Guest Editor: Mary McLeish

An Inquiry into Computer Understanding

Peter Cheeseman

NASA

Ames Research Center

This essay addresses a number of issues centered around the question of what is the best method for representing and reasoning about commonsense (sometimes called plausible inference), though Drew McDermott has shown that direct translation of commonsense reasoning into logical form leads to insurmountable difficulties from which he concluded that we must resort to procedural ad hoc-ery. This paper shows that the difficulties McDermott describes are a result of insisting on using logic as the language of commonsense reasoning. If instead, (Bayesian) probability is used, none of the technical difficulties found in using logic arise. For example, in probability, the problem of referential opacity cannot occur and non-monotonic logic (which, as McDermott shows, doesn't work anyway) is not necessary. The difficulties in applying logic to the real world are shown to arise from the limitations of truth semantics built into logic — probabilities substitute the more reasonable notions of belief. In Bayesian inference, many pieces of evidence are combined to get an overall measure of belief in a proposition. This is much closer to commonsense patterns of thought than long chains of logical inference through conclusions. Also, it is shown that English expressions of the "IF A THEN B" form are best interpreted as conditional probabilities rather than universally quantified expressions. Bayesian inference is applied to a simple example of linguistic information to illustrate the potential of this type of inference for AI. This example also shows how to deal with vague information; so far, that has been the province of fuzzy logic. It is further shown that Bayesian inference gives a theoretical inference for inductive inference that is born out in practice. Instead of insisting that probability is the best language for commonsense reasoning, a major point of this essay is to show that real inference is a complex interaction between probability, logic, and other formal representation and reasoning systems.

Respondents: Romas Aleliunas, Alan Bundy, N.C. Dalkey, A.P. Dempster, D. Dubois and H. Prade, Matthew Ginsberg, Russ Greiner, Pat Hayes, David Israel, Laveen Kanal and Donald Perlis, Henry Kyburg, Drew McDermott, D.L. McLeish, Charles Morgan, Eric Neufeld and David Poole, Judea Pearl, Larry Rendell, Enrique Ruspini, Len Schubert, Glenn Shafer, David Spiegelhalter, Ronald Yager, and Lofti Zadeh.

Abstracts for 4(2) May 1988 Special Issue on AI in France Guest Editor: Patrick Saint-Dizier

OBADE: Cognitive Modelling With Objects

Nelly Darcel

Université Paris XIII

Villetaneuse, France

Marie-Cl. Escarabajal

C.N.R.S.-Université Paris VIII

Saint Denis, France

We realize a computer simulation of children's reasoning in arithmetic word problem solving. The model parses the terms provided to the system in natural language and, while it performs this task, it tries to build its representation of the described situation by the way that a child elaborates a mental problem representation. This image results from three components: semantic knowledge, text comprehension process and problem-solving strategies.

We emphasize the adequacy between the knowledge representation and manipulation by an object formalism and the structure and the use of knowledge interacting in this application.

In our model the internal representation is realized in an object-oriented language whose main properties are accurately exploited. This choice allows us to combine the descriptive characteristics of each piece of knowledge with its implication in the progress of the process.

The program is supported by the analysis of individual protocols of some children: they allow us to hypothesize on the way the children modify their problem representation during the solving task.

We describe the main objects of the model. Then we simulate the terms of a problem, the way that the process is driven by expectations of contextually relevant information.

On the Consistency of Knowledge Bases: The COVADIS System

Marie-Christine Rousset

Université d'Orsay

Orsay, France

It is currently thought in the knowledge-based systems (KBS) domain that sophisticated tools are necessary for helping an expert with the difficult task of knowledge acquisition. The problem of detecting inconsistencies is especially crucial. The risk of inconsistencies increases with the size of the knowledge base; for large knowledge bases, detecting inconsistencies "by hand" or even by a superficial survey of the knowledge base is impossible. Indeed, most inconsistencies are due to the interaction between several rules via often deep deductions.

In this paper, we first state the problem and define our approach in the framework of classical logic. We then describe a complete method to prove the consistency (or the inconsistency) of knowledge bases that we have implemented in the COVADIS system.

On Fuzzy Syllogisms

Didier Dubois and Henri Prade
Université Paul Sabatier
Toulouse, France

This paper provides a systematic treatment of possibly imprecisely or vaguely specified numerical quantifiers in default syllogisms, following an approach initiated by Zadeh. The obtained propagation rules are derived from simple properties of relative cardinality, or equivalently (sic), conditional probability. Uncertainty in the description of numerical quantifiers is handled using probability theory, and particularly, fuzzy arithmetic. The advantages of this default reasoning method are its ability to model any kind of quantifier, and to build new defaults by chaining existing ones, in a rigorous manner. This approach also emphasizes the difference between two types of uncertain pieces of knowledge, i.e. conjectures versus general rules.

Towards An Analyzer (Parser) in a Machine Translation System Based on Ideas From Expert Systems

Zaharin Yusoff
Groupe d'Études pour la Traduction Automatique (GETA)
Saint-Martin d'heres, France

GETA is a research team working basically in the domain of machine translation. GETA's software system, ARIANE-78, has been tested over various pairs of relatively unrelated languages. Being a product of the late seventies, the system misses out on some of the artificial intelligence technology, particularly that of the eighties. Nevertheless, GETA carries out some research related to artificial intelligence within a general effort to bring improvements to the current system. This paper reports on an effort to embed such work within the framework of an entirely new system based on ideas from expert systems, significantly departing from the methodology of the current system (and that of other currently implemented machine translation systems). The proposed architecture aims for total modularity and flexibility, and some degree of intelligence.

A Question-Answering System for the French Yellow Pages

P. Herman, G. Sabah and A. Vilnat
LIMSI-CNRS
Orsay, France

This paper describes a dialogue-based system which is intended as an intelligent Natural Language interface to the French Yellow Pages.

We do not assume that the user knows how the Yellow Pages are organized, and we paraphrase his request, if necessary, so as to better search for the desired information. We do, however, assume that the reason the user is online is to find an address and phone number for some supplier.

There are three basic modules used in our system: Parser, Dialogue Manager and Generator. The first two exist (and are constantly being extended); the generation module is still only a set of functional specifications which will be outlined later in this article.

Processing of Unknown Words in a Natural Language Question-Answering System

J.P. Fournier, P. Herman, G. Sabah, A. Vilnat
LIMSI
Orsay, France
N. Burgaud, M. Gilloux
CNET
Lannion, France

This paper is an extension of "A Question-Answering System for the French Yellow Pages".

In this paper we first present a typology for the different kinds of unforeseen situations the system may have to handle while interacting with an untrained user. Then, we detail different strategies we have developed to solve those problems, more specifically to correct misspelled words and to deal with the properly typed, but unknown words. Those mechanisms have been integrated into the parser, and try to maintain, as long as possible, its determinism.

What Kind of Thing is a Concept?

Daniel Kayser
Université Paris-Nord

We discuss several features of concepts used for common knowledge: typicality, extensions of meaning, variability of type. We argue that these features are not superficial, lexical-level language-dependent issues, but deep characteristics of the knowledge itself. It is thus necessary to build knowledge representation systems compatible with these characteristics.

We show that many-valued and/or non-monotonic systems fail to capture typicality in its entirety. As for the two other features, no serious attempt has been made yet, and we only propose tentative elements for a solution.

The main idea is to decouple the notion of concept from the notion of basic element (predicate, node), and to represent a concept by an open-ended family of entities of the system. Each entity conveys a possible interpretation of the concept, and interpretations are ordered according to their "depth".

An example illustrating the main features of this scheme is provided.

Machine Learning Research at the Laboratoire de Recherche en Informatique at Orsay

Antoine P. Cornuejols
Université de Paris-Sud
Orsay, France

This note provides a brief account with sketchy details of the major directions in Machine Learning Research done at the Laboratoire de Recherche en Informatique (L.R.I.) at Orsay University in France. References contain certain publications giving details on the projects described below and closely related works.

Our research has several objectives: looking for a sound basis for the process of "generalization" from examples, using this to study conceptual clustering with automatic synthesis of descriptors, studying the nature and goodness of an "explanation" in the context of apprentice systems, and to develop experimental learning systems based on these principles applied to various practical domains. The approach taken by our research group has evolved with time but is still mainly

based on the learning of concepts from examples using logic representations and techniques. It corresponds to a major goal of our group: to give a clear and rigorous picture, if not a theory, of the topics under investigation. Of course, several tracks are being explored relating to different domains of tasks: learning of rule bases, games, computer aided teaching, learning in noisy environments, and so on. They are described below in the light of the main directions. The goal of a complete universal integrated system is still a far cry ahead, but as a famous Chinese proverb states, "The end lies in the way."

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, half the regular price, by using a copy of the all-purpose order form on the last page.

Please send article
submissions as
unformatted text

A.I. RESEARCH ANALYSTS VANCOUVER

MacDonald Dettwiler, a major international high-technology company with head office based in Richmond, B.C., has immediate career opportunities available for several A.I. professionals.

Since 1969 the company has delivered the world's highest quality computer based systems for the remote sensing and aerospace industries, focusing on resource management, image processing, meteorological, cartographic and flight control management applications.

In the A.I. field we have a dedicated group working on state-of-the-art projects relating to: knowledge-based image understanding, severe storm forecasting, space station maintenance task planning, intelligent equipment monitoring and advanced logical development shells.

If you are looking to join a growing and dynamic team committed to advances in A.I. technology we would like to hear from you.

A minimum of a Master's degree with a specialization in Artificial Intelligence is required along with directly related experience in applications oriented A.I. research.

Interested individuals can forward their resumes in confidence to:



**MACDONALD
DETTWILER**

3751 Shell Road, Richmond, B.C. V6X 2Z9
Telephone: (604) 278-3411 Fax: (604) 273-9830

ARE YOU WONDERING WHAT OPPORTUNITIES AWAIT YOU ?

Heuristic Search Inc. is Canada's first Personnel Agency dedicated to the Canadian Artificial Intelligence Industry.

CORPORATE CLIENTS

You can feel confident we understand your technical needs. If you are looking for seasoned AI Professionals for full or part time employment, contact us for further information.

AI PROFESSIONALS

AI Technologies will offer highly skilled computer professionals significant career opportunities that will command above average income levels for the near and long term future.

Call us for current opportunities, note there is no fee for this service.

**OUR SERVICE IS COMPLETELY CONFIDENTIAL AND
OUR GOAL IS YOUR SUCCESS !**

Please contact : Mr. Russell K. Senyk, P. Eng.
President

Heuristic Search Inc., 385 The West Mall, Suite 257, Toronto, M9C 1E7. Tel. 416-622-8129

International conference on Intelligent Tutoring Systems

Organized by: Université de Montréal

Sponsored by: CSCSI, AFCET, BCS

In Cooperation with: ACM -- SIGCUE and SIGART,
Inter-American Organisation for Higher Education

Conference Chairman:

Claude Frasson (Université de Montréal)

Program Committee Chairpersons:

Gregor V. Bochmann (Université de Montréal)

Marlene Jones (Alberta Research Council, Calgary)


Invited Speakers

Jacques Arsac (Université de Paris VI, France)
L'intelligence artificielle et l'enseignement assisté par ordinateur face au problème du sens
Jeffrey Bonar (LRDC, University of Pittsburgh, USA)
Intelligent Tutoring with Intermediate Representations
John Seely Brown (Xerox Park, Palo Alto, USA)
Steps Towards a New Epistemology for Situated Learning
John Self (University of Lancaster, UK)
Bypassing the Intractable Problem of Student Modelling
Elliot Soloway (Yale University, USA)
We Do Know More Now About Teaching and Learning
Pat Suppes (Stanford University, USA)
The future of Intelligent Tutoring Systems: Problems and Potential
Phil Winne (Simon Fraser University, Canada)
Representing and Inferencing with Knowledge about Teaching
Beverly Woolf (University of Massachusetts, USA)
Twenty Years in The Trenches: What Have We Learnt About Building Tutoring Systems
Masoud Yazdani (University of Exeter, UK)
Second Language Learning: An A.I. Approach

Panels

Stuart McMillan (SUN Microsystems, USA)
What can ITS do NOW?
Gordon McCalla (University of Saskatchewan, Canada)
What is wrong with ITS research?

Registration form

 **ITS 88** Complete and return to:
ITS 88
Claude Frasson
Département d'I.R.O.
Université de Montréal
C.P. 6128, Succ. "A"
Montréal, Canada H3C 3J7

Identification:

Name: _____ Given-name: _____

Title: _____ Organization: _____

Address: _____ City: _____

Province/State: _____ Country: _____

Postal-Code: _____ Phone: _____

I enclose my cheque for: ☐ The Conference
☐ The Banquet
☐ (____) Suppl. Proceeding(s)

Total (Can \$): _____

(Make cheques payable to ITS-88 Conference. No post-dated cheques.)

Montréal, June 1 - 3, 1988

Conference Objectives

ITS 88 will be a forum for presenting new results in research, development, and applications of ITS. The aims of the conference are to bring together specialists in the field of Artificial Intelligence and Education, to share state of the art information among attendees and to outline future developments of ITS and their applications. The Conference will include invited talks, panels and several sessions which will focus primarily on the following topics of interest:

- Learning environments
- Methodologies and architectures for educational systems
- AI programming environments for educational use
- Student modelling and cognitive diagnosis
- Curriculum and knowledge representation
- Knowledge acquisition in ITS
- Design issues in building ITS
- Practical use of ITS

Hotels facilities for the Conference

Participants are asked to directly reserve their rooms using the following list. These rooms will only be held until April 29th 1988, after which time availability cannot be guaranteed. Considering the numbers of rooms available at that period of the year, participants are advised to reserve as soon as possible.

Address	Phones	Rates
La Citadelle 410, Sherbrooke O. Montréal (Québec) Canada H3A 1B3	1-(514) 844-8851 1-(800) 361-1616 (CAN) 1-(800) 361-7545 (USA)	95.50 \$ /day (single) 95.50 \$ /day (double)
Shangrila 3407, Peel Montréal (Québec) Canada H3A 1W7	1-(514) 288-4141 1-(800) 361-7791 (CAN) 1-(800) 361-7791 (USA)	95.00 \$ /day (single) 105.00 \$ /day (double)
Château Champlain 1, Place du Canada Montréal (Québec) Canada H3B 4C9	1-(514) 878-9000 1-(800) 268-9420 (CAN) 1-(800) 828-7447 (USA)	112.00 \$ /day (single) 122.00 \$ /day (double)
Holiday Inn Centre-Ville 420, Sherbrooke O. Montréal (Québec) Canada H3A 1B4	1-(514) 842-6111 1-(800) HOLIDAY (World)	95.00 \$ /day (single) 115.00 \$ /day (double)

Conference fees (CAN \$):

	Before May 1st	After
ACM, CSCSI, AFCET, BCS members	250.00	275.00
Non members	275.00	300.00
Students (No Proceedings)	50.00	75.00
Banquet (June 2nd)	35.00	35.00
Proceeding (Suppl. copy)	50.00	50.00

Air Canada as official airline for the conference offers special convention fares to Canadian and American delegates by dialing this toll free number: 1-800-361-7585.
International delegates are encouraged to make early contact with any of Air Canada's offices or their travel agent for details on low cost, advance booking fares to Canada.

For further informations
please contact:

Gilles Gauthier
Phone: 1 (514) 282-8212

Email: gauthier%iro.udem.cdn@ubc

Upcoming Conferences

In Canada

**International Conference on
Intelligent Tutoring Systems**
1 - 3 June 1988, Montreal, Quebec
See announcement on page 54 for details.

**Artificial Intelligence '88
Vision Interface '88
Graphics Interface '88**
6 - 10 June 1988, Edmonton, Alberta
See announcements on pages 15 - 25 for details.

**First International Conference on Principles of
Knowledge Representation and Reasoning**
15 - 18 May 1989, Toronto, Ontario
See announcement on pages 36 - 37 for details.

**16th Conference of the Canadian
Association for Information Science**
12 - 14 May 1988, Ottawa, Ontario
Contact: David Holmes, c/o CAIS 88, PO Box 38, Station
B, Ottawa, Ontario K1P 6C3; (613) 564-4074

**Computer Processing of Chinese
and Oriental Languages**
29 August - 1 September 1988, Toronto, Ontario
Focus: Recent advances and development in
methodology, soft/hardware and expert knowledge.
Contact: Dr. Ching Y. Suen, Dept. of Comp. Sci.,
Concordia U., 1455 de Maisonneuve West, Montreal,
Quebec H3G 1M8.

**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.
Submission deadline: May 15, 1988. Submission material:
5 copies of an abstract (4 - 8 pages) or a full-length paper
(to 20 pages). There will be a best student paper award.
Contact: John Boose, Advanced Technology Center,
Boeing Computer Services, 7L-64 (via US mail:) PO Box
24346, Seattle, Washington, USA 98124; (206) 865-3253.
(via express mail:) Bldg 33.07 2760 160th Ave. SE
Bellevue, Washington, USA 98008. Brian Gaines, Dept. of
Comp. Sci., U. of Calgary, 2500 University Dr. NW,
Calgary, Alberta T2N 1N4 (403) 220-5901.

In the United States

**ICEBOL3 International Conference
on Symbolic and Logical Computing**
21 - 22 April 1988, Madison, South Dakota
Focus: ICEBOL is designed for teachers, scholars, and
programmers who want to meet to exchange ideas about
non-numeric computing. In addition to a focus on
SNOBOL, SPITBOL, and Icon, ICEBOL3 will feature
introductory and technical presentations on Prolog and
LISP, as well as on applications of BASIC, Pascal, and
FORTRAN for processing strings of characters. Topics
include: AI, expert systems, desk-top publishing, text
analysis.
Contact: Eric Johnson, 114 Beadle Hall, Dakota State
College, Madison, SD 57042; (605) 256-5270.
NET: eric @ sdnet (bitnet).

2nd Conference on Expert Database Systems
25 - 27 April 1988, Tysons Corner, Virginia
Topics include: Theory of knowledge bases; Object-
oriented systems; Reasoning on knowledge and
databases; Knowledge management; Distributed
knowledge and databases; Intelligent database interfaces;
Natural language interaction.
Contact: Larry Kerschberg, Dept. of Information
Systems and Systems Eng., George Mason U., 4400
University Drive, Fairfax, VA 22030.

**2nd Conference on Expert Systems and the Leading
Edge in Production Planning and Control**
3 - 5 May 1988, Charleston, South Carolina
Topics include: Process or product design; Production
control; Logistics; Scheduling; Integrated production
planning; Maintenance.
Contact: Michael Oliff, Dept of Management Science,
College of Business Administration, U. of South Carolina,
Columbia, SC 29208.

**AI and Advanced Computer Technology
Conference/Exhibition**
4 - 6 May 1988, Long Beach, California
The 24 technical sessions will be covered in 6 tracks
consisting of the following topics: Languages; Expert
Systems; Learning; Computers; Applications;
Engineering and Manufacturing.
Contact: Tower Conference Management Co., 331 W.
Wesley St., Wheaton, IL 60187.

8th Workshop on Distributed AI
22 - 25 May 1988, Lake Arrowhead, California
Topics include: Describing, decomposing, allocating
problems among a collection of intelligent agents;
Assuring coherent, coordinated interaction among
intelligent agents; Reasoning about other agents, the
world, and the state of the coordinated process;
Recognizing and resolving disparities in viewpoints,
representations, knowledge, goals, etc.; Problems of

language and communication.

Contact: Les Gasser, Distributed AI Group, Comp. Sci. Dept., U. of Southern California, Los Angeles, CA. 90089-0782; (213) 743-7794.

9th Conference on Automated Deduction

23 - 26 May 1988, Argonne, Illinois (near Chicago)

Topics include: Theorem-proving; Logic programming; Unification; Deductive databases; Term rewriting; Automatic theorem-proving for non-standard logics; Program verification; Inference systems.

Contact: Ewing Lusk, CADE-9, Mathematics and Comp. Sci. Division, Argonne National Laboratory, Argonne, IL 60439.

Illinois Interdisciplinary Workshop on Decision Making

15 - 17 June 1988, Champaign-Urbana, Illinois

Focus: Representation and use of knowledge for decision making in human, mechanized, and ideal agents, from the fields of AI, philosophy, psychology, statistics, and operations research. Topics include: The representation, organization and dynamics of the knowledge used in decision making; Decision making strategies; Decisions under constraints (limited rationality); Combining normative and descriptive theories; The use of domain knowledge to initialize beliefs and preferences.

Contact: L. Rendell, Dept. of Comp. Sci., U. of Illinois, 1304 W. Springfield Ave., Urbana, IL 61801.

1st Conference on Industrial and Engineering Applications of AI and ES

2 - 3 June 1988, Tullahoma, Tennessee

Topics include: Fault monitoring and diagnosis; Automatic controls; CAD/CAM; Software engineering; Systems engineering; Medical and biomedical engineering; Aerospace engineering; Oceanic and atmospheric engineering; Chemical, petroleum and mining engineering; Human-machine interfaces; Space station automation; Robotic vision; Aircraft and flight management; Operations in nuclear or hazardous environments; Command control and communications.

Contact: Moonis Ali, U. of Tennessee Space Institute, Tullahoma, TN 37388.

26th Annual Meeting of the Association for Computational Linguistics

7 - 10 June 1988, Buffalo, New York

Topics include: Computational linguistics; Phonetics, phonology, and morphology; Interpreting and generating spoken and written language; Models of language; Machine translation aids; Natural language interfaces; Message understanding systems.

Contact: Jerry Hobbs, ACL88 Program Chair, AI Center, SRI International, 333 Ravenswood Ave., Menlo Park, CA 94025; (415) 859-2229.

NET: hobbs@warbucks.ai.sri.com.

Symposium on the Engineering of Computer-Based Medical Systems

8 - 10 June 1988, Minneapolis, Minnesota

Topics include: Expert Systems; Networking and communication; Software QA; Reliability; Regulatory

issues; Medical imaging and graphics.

Contact: Bart Galle, Continuing Medical Education, Box 202 UMHC, 420 Delaware St. SE, Minneapolis, MN 55455; (612) 626-5525.

5th Conference on Machine Learning

12 - 15 June 1988, Ann Arbor, Michigan

Topics include: Concept learning, Genetic algorithms, Empirical evaluation, Learning and planning, Empirical methods, Theoretical analysis, Language learning, Psychological validity, Learning and design, Connectionist, Machine discovery, Analogical reasoning.

Contact: Machine Learning Conference, Cognitive Science and Machine Intelligence Laboratory, U. of Michigan, 904 Monroe St., Ann Arbor, MI 48109-1234.

Internet: ml88@csmil.umich.edu.

2nd Conference on Theoretical and Methodological Issues in Machine Translation of Natural Language

12 - 14 June 1988, Pittsburgh, Pennsylvania

Topics include: Machine-aided translation; Automatic analysis and generation of natural language texts; Structure of lexicons and grammars; Research tools and Methodologies; Theory of translation.

Contact: Cerise Josephs, Center for Machine Translation, Carnegie-Mellon U., Pittsburgh, PA 15213; (412) 268-6591. NET: cerise@nl.cs.cmu.edu.ARPA.

1988 American Controls Conference

15 - 17 June 1988, Atlanta, Georgia

Focus: Neural applications to robotics.

Contact: Wayne Book, The George W. Woodruff School of Mechanical Engineering, Georgia Institute of Technology, Atlanta, Georgia 30332 (404) 894-3247.

COMPASS '88

3rd Annual Computer Assurance Conference

27 June - 1 July 1988, Washington, D.C.

Focus: Explore computer assurance, define its boundaries, identify its issues, and realize its objectives.

Contact: Frank Houston, P.O. Box 5314, Rockville, MD 20851

NET: houston@nrl-csr.arpa

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 (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, Rice U. Dept of Comp. Sci.,
P. O. Box 1892, Houston, TX 77251-1892; (713) 527-4834.
NET: cork@rice.edu.

ASME Computers in Engineering Conference

31 July - 3 August 1988, San Francisco, California

Topics include: CAD/CAM; Computer simulation; AI; Robotics; Interactive graphics; Finite element techniques; Microprocessors; Computers in education.
Contact: Edward Patton, US Army Ballistic Research Lab, Aberdeen Proving Grounds, MD 21005; (301) 278-6805.

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 student participants.

Participation deadline: April 15. Send a description of current research, 1 page max. Participation is limited to 60. Submission deadline: April 15. Submission material: 8 copies of an extended abstract of max. 5 pages, with 200 word summary.

Contact: John Cherniavsky, Workshop on Computational Learning Theory, Dept of Comp. Sci., Georgetown U. Washington, D.C. 20057

3rd Conference on Applications of AI in Engineering

8-12 August 1988, Stanford U., California

Topics include: Representation in design; Problem solving in diagnosis/evaluation; Constraint reasoning in process control and planning; Robotics learning; Tutoring qualitative models; User interfaces.

Contact: John Gero, Technical Chair, AIE88, Dept. of Architectural Science, The U. of Sydney, NSW 2006 Australia; Phone: 61-2-692-2328.

UUCP: uunet!munni!archsci.su.oz!john.

3rd International Conference on CAD/CAM Automation Robotics and Factories of the Future

14 - 17 August 1988, Southfield, Michigan

Topics include: Manufacturing workcell diagnosis; Process planning; Robot motion planning; Scheduling; Knowledge representation of workcells; Sensor-based programming; Vision; Object representation.

Contact: Biren Prasad, Electronic Data Systems, EDS Pinehurst #201, 1400 North Woodward Ave, Bloomfield Hills, Michigan 48013 (313) 645-4714.

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; (714) 856-5517.

Sponsored by Computer Professionals for Social Responsibility, P.O. Box 717, Palo Alto, CA 94301.

AAAI 4th Workshop on Uncertainty in AI

19 - 21 August 1988, St. Paul, Minnesota

Topics include: Applications: results, implementation problems and experiences, analyses of the experiences of end users; Knowledge engineering under uncertainty; Control of uncertain reasoning processes; Different uncertainty calculi; Robotics; Planning.

Contact: Ross Shachter, Center for Health Policy, 125 Old Chemistry Building, Duke U., Durham, NC 27706; (919) 684-4424.

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

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.

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.

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 ES; Machine learning, cognition and the cockpit; Neural networks and human-machine interfaces; Parallel processing and NN; 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 Conference on Computer Vision

5 - 8 December 1988, Tarpon Springs, Florida

Focus: All aspects of computer vision.

Submission deadline: May 15. Submission materials: 4 copies of complete drafts.

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

2nd International Workshop on AI and Statistics

4 - 7 January 1989, Ft. Lauderdale, Florida

Topics include: Consultation systems; Statistician's assistants; Knowledge representation for statistics; Uncertainty propagation; Clustering and concept formation; Statistical methods for knowledge acquisition. Contact: W. Gale, AT&T Bell Laboratories, 2C278, 600 Mountain Ave., Murray Hill NJ 07974.

IFIP Congress '89: Better Tools for Professionals

28 August - 1 Sept 1989, San Francisco, California

Program Tracks: Fundamental tools; Languages and operating systems; Communication and distributed systems; Knowledge-based systems; Software engineering; Supercomputing; VLSI-CAD tools; Office automation; Factory automation; Education; Computers and society.

Submission deadline: Nov 1, 1988. Submission material: 6 copies of final form papers including abstract, max. 4500 words.

Contact: Herve Gallaire, ECRC, Arabellastr 17, D-8000 Numich 81 FRG. Phone: 49-89 92 69 91 00.

Outside North America

International Workshop on Microcomputers and Expert Systems

26 - 28 April 1988, Hong Kong

Focus: To describe the development of over 15 microcomputer-based packages expressly designed for planners, engineers and management. Expert systems applications under development will be presented.

Contact: Ron Sharpe, Commonwealth scientific and Industrial Research Organisation (CSIRO), Div. of Construction and Engineering, PO Box 56, Highett, Vic 3190, Australia. Phone (+61 3) 556 2211, Fax (+61 3) 553 2819, Telex AA33766 e-mail: ron@dbrmelb.dbrhi.oz.

Avignon 88-Expert Systems and Applications

30 May - 3 June 1988, Avignon, France

The main conference will be accompanied by specialized conferences: Expert Systems and Maintenance; Expert Systems in Medicine and Biology; AI and Defense. Topics include: Finance, banking, economics; Business administration, office automation; CAM, robotics, QA; CAD, CAT, engineering; Security, software engineering. Contact: Jean-Claude Rault EC2 269-287, rue de la Garenne, 92000, Nanterre, France.

Phone: (33.1) 47 80 70 00.

nEURO'88 1st European Conference on Neural Networks

6 - 9 June 1988, Paris, France

Topics include: Models of memory and learning, sensory perception, motor control; Methods for solving specific problems with artificial neural networks; Artificial network architectures, electronic and optical implementations, and applications including robotics.

Contact: nEURO'88, G. Dreyfus E.S.P.C.I., 10 rue Vauquelin, F-75005 Paris, France.

ACM-SIGIR

11th Conference on Research and Development in Information Retrieval

13 - 15 June 1988, Grenoble, France

Topics include: Retrieval system modeling; Retrieval and AI; Evaluation techniques; Natural language processing; Database management; User interfaces.

Contact: Gerard Salton, Cornell U. Dept. of Comp. Sci., 4130 Upson Hall, Ithaca, NY 14853 - 7501, USA. NET: siri@imag.UUCP.

EKA88 2nd European Knowledge Acquisition for Knowledge-based Systems Workshop

19 - 23 June 1988, Bonn, Germany

The limited attendance of 40 will be confined to those presenting their work. Topics include: Transfer/modeling of expertise; Apprenticeship, explanation-based, and other learning systems; Issues in cognition and expertise that affect the knowledge acquisition (KA) process; Extracting and modeling of knowledge from text; Integration of KA techniques within a single system; KA methodology and training; Validation of KA techniques.

Contact: Brian Gaines, Dept. of Comp. Sci., U. of Calgary, 2500 University Dr. NW, Calgary, Alberta, Canada T2N 1N4 (403) 220-5901.

Role of AI in Databases and Information Systems

4 - 8 July 1988, Guangzhou, China

Focus: The Role of AI in Databases and Information Systems and the role of Databases in AI based systems. Features: John McCarthy, Stanford U. "Knowledge about knowledge in databases"; John Sowa, IBM "Knowledge representation in databases, information systems and natural language"; Raymond Reiter, U. of Toronto "Integrity constraints in databases and knowledge bases". Contact: A. Solvberg, Norwegian Inst. Techn., U. Trondheim, Norway.

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, Università degli Studi di Milano, Via Moretto da Brescia, 9 20133 Milano, ITALY. Phone: +39-2-2141230.

PANEL '88 EXPODATA 17 JAIIO XIV Latin American Conference on Information Sciences

September 1988 Buenos Aires, Argentina

Topics include: Databases; Distributed systems; Knowledge based systems; Expert systems; Theory of computation; Robotics; Computer architecture; CAM; Software engineering; Functional and logic programming; New applications in education, management, health, law, social sciences, linear programming, optimization.

Submission deadline: June 15, 1988. Submission

material: max 20 pages with 200 hundred word abstract.
 Contact: PANEL '88/17 JAIIO, SADIO Uruguay 252/2do
 Piso "D", 1015 Buenos Aires, ARGENTINA. Phone: 54 (1)
 40-5755 or 45-3950

**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.

Submission deadline: May 1, 1988. Submission material: 5 copies of papers, 5000 words max. with 200 word abstract.

Contact: Derek Sleeman, Dept of Comp. Sci., U. of Aberdeen., ABERDEEN AB9 2UB, Scotland UK. Phone: Aberdeen (+44 224) 272288. Telex 73458.

**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.

Submission deadline: May 15, 1988. Submission material: 2 copies of a 400 word summary.

Contact: Mikio Takagi, Institute of Industrial Science, U. of Tokyo, 7-22-1, Roppongi, Minato-ku, Tokyo 106.
 CSNET: takagi%iiis.u-tokyo.junet@relay.cs.net.

1st Australian Knowledge Engineering Congress

15 - 17 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 KE; Design of Intelligent tutors.

Submission deadline: "ASAP". Submission material: A preliminary indication of interest in offering a paper, presentations.

Contact: G. Garner, Deakin U., Victoria 3217, Australia.

NET: brian@aragorn.oz.

**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.

Submission deadline: June 15. Submission material: 4 copies of a paper max. 5000 words with 100 word abstract.

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 areas: Finance, banking, insurance, economics, DSS, public and private services, OA, 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.

Advertiser Index

Addison - Wesley	38	KR Conference	36 - 37
Applied AI	Cover IV, 6	Lisp Canada	30 - 31
Autometrics	44	MacDonald Dettwiler Associates	53
CAIP Corporation	Cover III	Neuron Data	5
Gold Hill	Cover II	University of Toronto	43
Heuristic Search	53	University of Waterloo	43
ITS Conference	54		

Advertising Notes. Those interested in advertising in the magazine, please write us to obtain a Press Kit. Advertisers who reserve space for four consecutive issues are eligible for discounted rates.

Use this form to join CSCSI/SCEIO, to subscribe to
Canadian AI or Computational Intelligence, and to order publications

Canadian Society for
Computational Studies
of Intelligence

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

Application for Membership and / or Journal and Conference Proceedings Order

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

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

- ☐ I wish to join CSCSI/SCEIO and receive *Canadian AI* (\$35).
☐ I am a student (subtract \$10). **AND/OR** ☐ I am a member of CIPS (subtract \$10).
- ☐ Please send me the following back issues of *Canadian Artificial Intelligence* (\$10 each including postage, except \$15 for #3):*.....
- ☐ Please enter my subscription to *Computational Intelligence* at the CSCSI/SCEIO non-institutional member discount rate (\$16).†
- ☐ Please send me the following CSCSI/SCEIO conference proceedings (\$30 for 1986; \$25 for 1982 and 1980; 1984 is out of print. Add \$5 for postage within Canada, \$7 outside):
☐ Montréal, 1986. ☐ Saskatoon, 1982. ☐ Victoria, 1980.

Total enclosed: \$..... Cdn/US (Payment may be made in U.S. dollars at the current rate of exchange.)

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



DELIVERING EXPERT SYSTEMS

LEADERSHIP

CAIP Corporation is the leading Canada-based organization concentrating exclusively on practical applications of Artificial Intelligence technology. Since its inception in 1984, CAIP has developed extensive experience in the analysis and development of knowledge-based expert systems, the fastest growing segment of AI.

RESOURCES

CAIP Professionals comprise experienced knowledge engineers, business consultants, systems analysts and many of Canada's leading scientists from universities across the country. The ability to configure project teams that include world renowned academics provides CAIP with a unique service.

EXPERIENCE

CAIP is a member of International Artificial Intelligence. This affiliation provides CAIP with an unparalleled synergy with many of the leading companies in the fields of AI workstation hardware, AI software shells, mainframe database management systems and office systems technology.

SERVICES

CAIP provides a full range of client services to organizations throughout Canada.

- EXPERT SYSTEM DEVELOPMENT
- APPLICATION SURVEYS
- FEASIBILITY STUDIES
- PROTOTYPE DEVELOPMENT
- CONTRACT RESEARCH AND DEVELOPMENT
- EDUCATION
- PRODUCT MANAGEMENT

106 Colonnade Road, Suite 220, Ottawa, Ontario K2E 7P4
(613) 727-0082

5915 Airport Road, Suite 200, Toronto, Ontario L4V 1T1
(416) 671-4650



CAIP CORPORATION

Canadian Artificial Intelligence Products

A member of the NEXA Group.

System Development and Prototyping

Prototype/Pilot Systems Developed:

- Knowledge-Based Transit Dispatcher's Assistant
- Spacecraft Autonomy Management
- Knowledge-Based Database Protocol Conversion
- several other small scale in-house demonstrators

Presently Under Development:

- Expert System for Highway Transportation Dispatch
- Knowledge-Based Proposal Evaluation Assistant

Feasibility Studies Conducted:

- Application of AI Techniques in Spacecraft Operation Monitoring
- Fault-Tolerant Computing and Artificial Intelligence
- Application of Expert System Technology in Canadian Transportation (jointly with Cognos Incorporated)
- Application of Expert Systems in Marine Transportation
- Residency Expert System Study

Studies Underway

- Application of Voice Recognition/Synthesis in Transportation
- Schanckean Natural Language Understanding
- Neural-Network Systems

AI System Development Tools

"Q&A"

Database/word processing package with an extensive Natural Language front end. **SYMANTEC** 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, \$445.


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, \$550. Time Line Graphics—an 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. Available for use on IBM AT, 386, PS/2 and Macintosh II. \$6650.



ACCENT SA

A high quality text-to-speech converter 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, \$805.

AICOM CORPORATION

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 interpreter/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, \$1507.



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.

Expertech.

ENQUIRY AND ORDER: Call (613) 592-0084 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.



Applied AI Systems, Inc.
4015 Carling Ave,
KANATA, Ontario, Canada K2K 2A3
Telephone: (613) 592-0084
Facsimile: (613) 592-2333

LPA PROLOG Professional, LPA MacPROLOG and Apes are registered trademarks of Logic Programming Associates Ltd. Q&A is a registered trademark of Symantec Inc. Time Line is a registered trademark of Breakthrough Software. ACCENT SA is a registered trademark of AICOM CORPORATION. NeuralWorks Professional is a registered trademark of NeuralWare Inc. Nexpert is a registered trademark of Neuron Data. Xi Plus is a registered trademark of Expertech.