Canadian Artificial Intelligence
Intelligence Artificielle au Canada

Winter 1995 | No. 36 | hiver 1995

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 or

IJCAI: International Jamboree for Canadian Artificial Intelligence
IJCAI: Foire Internationale pour l'Intelligence Artificielle Canadienne

AI Summer in Montreal
Éte d'IA à Montréal

Comdale: A Canadian AI Success Story
Comdale: Un Récit de Succès d'IA au Canada

Pure and Applied Research in the Biological Co
Recherche Pure et Appliquée dans le Projet de

Artificial Intelligence and the World Wide Web
L'Intelligence Artificielle et le World Wide Web

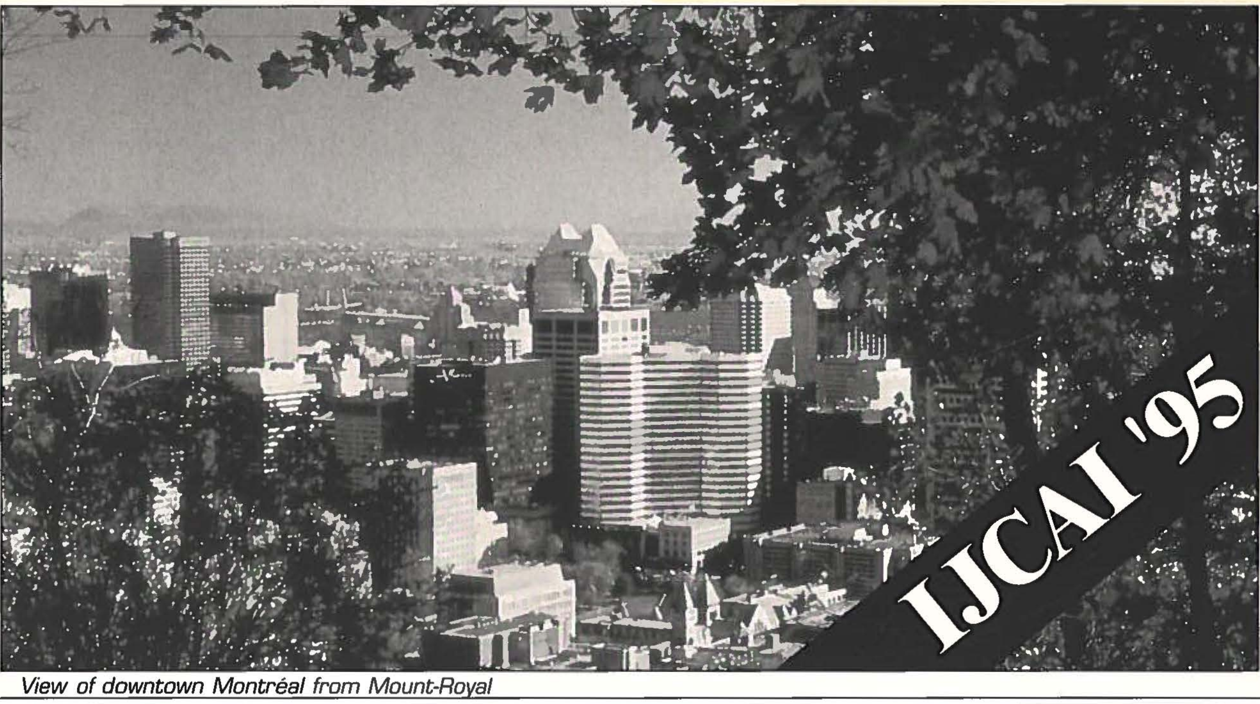

Photo Credit: Greater Montréal Convention and Tourism Bureau

View of downtown Montréal from Mount-Royal

Retrospective

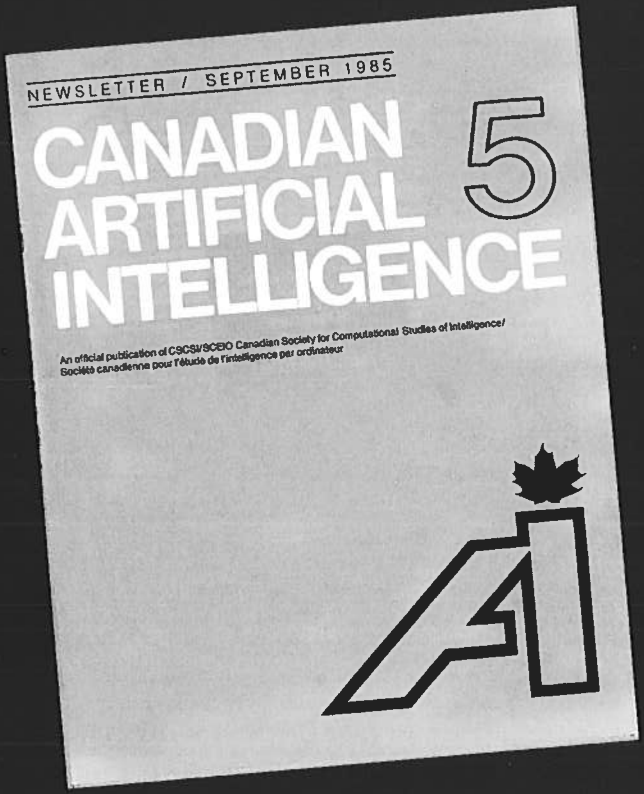

Contents

Editor rebuts self in exclusive interview (*Graeme Hirst*), 5
New bindings, 6
Canada is prominent in IJCAI awards, 8
Research funding, 10
Company news, 11
Cartoon (*P.S. Mueller*), 11
Xerox announces new low-cost AI workstations, 14
Workshop on Theoretical Approaches to Natural Language Understanding (*Graeme Hirst* and *Fred Popowich*), 15
Workshop on Foundations of Adaptive Information Processing (*Vijay Raghavan*), 15
Canada conquers Los Angeles (*Gord McCalla*), 20
Directory of Canadian AI companies, 25
New books and journals, 26
Obituaries, 30
Recent technical reports, 32
Cartoon (*Johnny Ng*), 37
Forthcoming conferences, and calls for papers, 39
All-purpose form, 43

Cover graphic, contents from Issue Number 5, September 1985 Canadian Artificial Intelligence Newsletter

Canadian Artificial Intelligence

Intelligence Artificielle au Canada

Winter 1995 | No. 36 | hiver 1995

Canada's National AI magazine.

Editors · Editeurs
Peter Turney and Sue Abu-Hakima

Editor Emeritus · Rédacteur emeritus
Roy Masrani

Managing Editor · Editeur en chef
Arlene Merling

Production Team · Equipe de production
Arlene Merling, Greg Klymchuk, Carol Tubman

Academia · Académiques
Vacant

Canadian AI Success Stories
Récite de Succès d' IA au Canada
Peter Turney

PRECARN Update · Nouvelles de PRECARN
Jean-Claude Gavrel

Book Reviews · Critiques de livres
Graeme Hirst

Translation · Traduction
Michel Addison, Chantal Desormeaux-Preston, Benoit Farley

Advertising · Publicité
Arlene Merling

Canadian Artificial Intelligence is published three times a year 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). Canadian Publications Mail Product Sales Agreement No. 507032.

ISSN 0823-9339

Copyright © 1995, 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 Tri-ad Press Ltd. *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, the National Research Council or the Alberta Research Council.

Copyright © 1995, 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 Tri-ad Press Ltd. *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, de Conseil National de Recherches du Canada, ou du Conseil de Recherche de l'Alberta.

CSCSI Executive

President · Président:
Stan Matwin, Computer Science Dept., University of Ottawa, Ottawa, ON K1N 6N5; stan@ csi.uottawa.ca

Past-President · Président Précédent:
Janice Glasgow, Dept. of Computing & Information Science, Queen's U., Kingston, ON K7L 3N6; janice@qucis.queensu.ca

Vice-President · Vice-Président:
Allan Jepson, Department of Computer Science, University of Toronto, Toronto ON M5S 1A4; jepson@vis.utoronto.ca

Secretary · Secrétaire:
Fred Popowich, School of Computing Science, Simon Fraser University, Burnaby BC V5A 1S6; popowich@cs.sfu.ca

Treasurer · Trésorier:
Peter van Beek, Dept. of Computing Science, U of A, Edmonton, AB T6G 2H1; vanbeek@cs.ualberta.ca

Editors · Editeurs:
Peter Turney, Knowledge Systems Laboratory, NRC, Ottawa, ON K1A 0R6; peter@ai.iit.nrc.ca
Sue Abu-Hakima, Knowledge Systems Laboratory, NRC, Ottawa, ON K1A 0R6; sue@ai.iit.nrc.ca

Contents | Contenu

Contents	Page	Contenu
Communications	2	**Communications**
Editorial Artificial Intelligence and the World Wide Web *Peter Turney*	4	**Editorial** L'Intelligence Artificielle et le World WideWeb *Peter Turney*
Feature Articles IJCAI: International Jamboree for Canadian Artificial Intelligence *Alan K. Mackworth*	7	**Gros Titres** IJCAI: Foire Internationale pour l'Intelligence Artificielle Canadienne *Alan K. Mackworth*
AI Summer in Montreal *C. Raymond Perrault*	9	Éte d'IA à Montréal *C. Raymond Perrault*
Canadian AI Success Stories Comdale: A Canadian AI Success Story *Aurora C. Diaz*	10	**Récite de Succès d' IA au Canada** Comdale: Un Récit de Succès d'IA au Canada *Aurora C. Diaz*
Academia Pure and Applied Research in the Biological Computation Project *Michael R.W. Dawson*	17	**Académiques** Recherche Pure et Appliquée dans le Projet de Calcul Biologique *Michael R.W. Dawson*
PRECARN Update	23	**Nouvelles de PRECARN**
Book Reviews	24	**Critiques de livres**

Canadian Artificial Intelligence welcomes submissions on any matter related to artificial intelligence. Please send your contribution, electronic preferred, with an abstract, a photograph, and a short bio to:

Dr. Peter Turney
Co-Editor, *Canadian Artificial Intelligence*
Knowledge Systems Laboratory, Institute for Information Technology
National Research Council Canada
M-50 Montreal Road
Ottawa, Ontario, Canada
K1A 0R6 or — peter@ai.iit.nrc.ca
Telephone (613) 993-8564 Fax (613) 952-7151

Advertising rates are available upon request.

Arlene Merling
Managing Editor, *Canadian Artificial Intelligence*
Alberta Research Council
3rd Floor, 6815 - 8 Street N.E.
Calgary, Alberta, Canada
T2E 7H7 or — merling @arc.ab.ca
Telephone (403) 297-2608 Fax (403) 297-2505

Book reviews and candidate books to review should be sent to:

Graeme Hirst
Canadian Artificial Intelligence
Department of Computer Science
University of Toronto
Toronto, Ontario, Canada
M5S 1A4 or — gh@ai.toronto.edu

Recycled / Recyclable

COMMUNICATIONS
COMMUNICATIONS

Treasurer's Report

Canadian Society for Computational Studies of Intelligence (CSCSI) Treasurer's Annual Report for the 1994 Fiscal Year

CSCSI had two major financial activities in the 1994 fiscal year — the AI'94 conference held in Banff, Alberta and the production of the magazine. A financial report for the conference follows this report. The profits from the conference helped the society to almost break even in 1994.

Income	
Membership Dues	$ 8, 162.03
Magazine Advertising	
CAI 33, Winter '94	747.67
CAI 34 Spring/Summer '94	800.00
CAI 35, Autumn '94	660.00
GST Collected	906.74
GST Refund from Government	46.62
Interest on Bank Account and GIC	141.40
AI'94 Conference and Workshop Income	4, 924.19
Total Income	**$ 16, 388.65**
Expenses	
Administration Fee for Collecting Dues	1, 890.00
Magazine Production	
CAI 33, Winter '94	5, 730.99
CAI 34, Spring/Summer '94	4, 259.05
CAI 35, Autumn '94	5, 492.53
GST paid out	648.49
GST paid to government	304.87
Total Expenses	**$ 18, 325.93**
Year Total	**$ (1, 937.28)**
Bank Balance as of 31/12/94 (incl. GIC's)	**$ 31, 847. 69**

Notes

I have decided to present the magazine income and expenses according to when the issue was published, rather than in what year the money was actually received or spent. Some of the advertising income has not yet been received. The Canadian Information Processing Society (CIPS) handles our membership renewals and charges us a fee of $ 4.00 for renewals who are also members of CIPS and $10.00 for renewals who are not members of CIPS. In 1994, there were 246 renewals (95 were members of CIPS and 151 were not members of CIPS).

Treasurer's report for the AI'94 conference

CSCSI held its 10th conference May 16-20, 1994 in Banff, Alberta. Artificial Intelligence '94 (AI'94) was held jointly with Graphics Interface '94 and Vision Interface '94 under the title AI/GI/VI'94. Speaking as the treasurer, the conference was a success from the financial point of view and, speaking personally, the conference was a success from the intellectual and social points of view (I attended some interesting talks and got to talk to some people I hadn't seen since the last conference). For AI'94, two financial statements have been produced: one that covers income and expenses for the conference as a whole and another that considers income and expenses particular to AI'94. Below I show only the financial picture particular to AI'94.

Income	
Conference grant, NSERC	$ 3, 505.00
Conference grant, U. of Alberta	3, 000.00
Page charges (extra pages in proceedings)	244.82
AI proceedings used at the conference	3,150.00
Workshop registrations	1,130.06
CSCSI's share of profits from conference	4,826.43
Total Income	**$ 15,856.31**
Expenses	
Invited speaker expenses	6, 656.57
Workshop proceedings	673.19
Conference proceedings	3, 097.43
Postage and couriers (program chair)	504. 93
Total Expenses	**$ 10, 932.12**
Conference Total	**$ 4, 924.19**

Notes

The AI/GI/VI'94 conference paid CSCSI for every AI proceedings it gave out as part of a registration.

Two AI workshops were organized and all of the profits for these workshops were given to CSCSI. Denis Gagne of College Militaire Royal organized the Distributed AI

Workshop and Bruce MacDonald of the University of Calgary organized the Machine Learning Workshop. Thanks once again to both of them for their efforts.

Considering just the shared income and expenses for the conference as a whole, the AI/GI/VI'94 conference made a profit of $ 12,066.07. CSCSI's share was two-fifths which is based on the proportion of overall registrants that attended the AI portion of the conference.

Peter van Beek
Treasurer, CSCSI

New Incentive for Students to Join CSCSI

There is good news for our regular subscribers — CSCSI membership rates remain unchanged this year. There is even better news for student subscribers — rates have been slashed in half to encourage you to take advantage of the benefits of CSCSI membership. Student memberships now cost a mere $ 15.00. Institutional membership rates have increased to reflect the wider distribution of a single copy.

A major benefit of membership in CSCSI is *Canadian Artificial Intelligence*. It was started in 1974 as the CSCSI/SCEIO Newsletter, and was reborn in 1984 through the sterling efforts of Professor Graeme Hirst of the University of Toronto. Recent issues have been produced in Calgary with generous assistance from the Alberta Research Council, and the tradition of quality has been upheld. The magazine contains articles of general interest, descriptions of current research and courses, reports of recent conferences and workshops, news about AI companies and products, paid advertising, announcements of forthcoming activities, etc. Sometimes you can even find humour and cartoons!

Members also receive reduced rates for:

- CSCSI Conference Registration
- CSCSI Conference Tutorials
- Subscription to Computational Intelligence
- Subscription to Machine Learning

Please see the inside back cover for the application form.

IJCAI Awards for 1995

On behalf of the IJCAI trustees, I'm happy to announce that Herb Simon is the winner of the IJCAI Research Excellence Award for 1995. The Computers and Thought award will be shared by Stuart Russell and Sarit Kraus.

Ray Perrault
General Chair
IJCAI '95
perrault@ai.sri.com

Errata

Gremlins were at work in our last issue of Canadian Artificial Intelligence magazine. The original figure which was to accompany Lawrence Hunter's article "Artificial Intelligence and Molecular Biology 1994" was submitted in colour. Our attempts to re-create the graphic to differentiate between the red and green areas using gray scale levels was unsuccessful. The original figure, "A comparison of the predictions of Reinetz, et al.'s computer model of gene expression with actual expression levels, using a modified representation of the original images," can be viewed on the web (see http://ai.iit.nrc.ca/cscsi/cai/hunter.html).

Our sincere apologies to Lawrence, and to our readers for the resulting confusion.

Peter Turney

EDITORIAL
EDITORIAL

Artificial Intelligence and the World Wide Web

Introduction

Qu'est-ce que l'Internet a à offrir aux travailleurs d'Intelligence Artificielle? Un accès instantané à des rapports techniques provenant du monde de l'informatique, qui prennent la forme de fichiers PostScript; des mégaoctets de logiciels gratuits en Intelligence Artificielle, incluant le code source; des "pages WWW personnelles," ou des personnes travaillant en IA decrivent leurs interets de recherche personnelles et produisent une liste de leur travaux de recherche (fréquemment disponibles sous forme de fichiers PostScript) et (occasionnellement) fournissent une photographie; des revues électroniques, quelquers-unes en format hypertexte, la plupart offrant une publication très rapide avec de très grands standards; solicitation de travaux de recherche pour des conférences d'IA futures; bibliographies interactives de littératures d'IA; demonstration en-ligne de logiciel d'IA (rechercher le démo Fuzzy Shower de Knowledge Systems Lab); et plus encore.

Introduction

What does the Internet offer for AI workers? Instant access to thousands of computer science technical reports, in the form of PostScript files; megabytes of free AI software, including source code; "personal home pages," where people working in AI describe their personal research interests and (typically) list their papers (frequently available as PostScript files) and (occasionally) include a photograph; on-line electronic journals, some in hypertext format, most offering very rapid publication and high standards; CFPs (Calls For Papers) for all upcoming AI conferences; searchable bibliographies of AI literature; on-line demonstrations of AI software (look for the Knowledge Systems Lab's Fuzzy Shower demo); and more.

Many readers of Canadian AI are already familiar with the World Wide Web and its value to those who work in the field of Artifical Intelligence. However, some of our readers are not (yet) on the Internet and others who are on the Internet do not (yet) know how to take advantage of this resource. This article begins with a brief introduction to the Internet and the World Wide Web (aka "the web") and then discusses web resources that can be helpful for AI workers. For those who do not yet have Mosaic (a favourite tool for browsing the web — other popular tools are Lynx and Netscape), we explain how to get it. For those who do not yet have Internet access, we explain how to get that too.

The Internet

The Internet is the world's largest computer network and it is currently in a phase of rapid expansion. There are between two million and three million registered Internet addresses with an estimated average of ten people per address. At 20 to 30 million people, the Internet is about 20 times larger than, for example, Compuserve.

The Internet is not a single network governed by a single group. It is a large number of separate networks that have organically fused together. What unites these separate networks, besides physical connections, is a common set of communication protocols, called the Internet Protocol Suite (IPS). This set of protocols defines a standard for each type of traffic. For example, SMTP (Simple Mail Transfer Protocol) is the standard for electronic mail and FTP (File Transfer Protocol) is the standard for exchanging files. The IPS is continuously evolving. Anyone can propose a new standard, by writing and circulating a document called a RFC (Request For Comments). In 1989, a new protocol was proposed by a group of physicists at CERN (Conseil Européen pour la Recherche Nucléaire, in Geneva, Switzerland). The protocol was called HTTP (HyperText Transfer Protocol) and was designed to facilitate the exchange of hypertext over the Internet. The part of the Internet traffic that uses HTTP is called the World Wide Web.

The World Wide Web

Like FTP, HTTP is a client-server protocol. There are HTTP servers, which supply or transmit hypertext, and there are clients, which browse or receive hypertext. The first HTTP client software, developed at CERN, had a simple text-based interface. In 1993, when NCSA (National Center for Supercomputer Applications, at the University of Illinois) developed a graphical HTTP client, called Mosaic, the web began a serious growth spurt. There are now more than a million HTTP clients and more than 8,000 HTTP servers on the Internet. HTTP has grown to become the third most common form of traffic on the Internet. (Number one is FTP and number two is NNTP. NNTP is the protocol used by news groups to exchange news messages.)

Mosaic is essentially a point-and-click interface to the Internet. Mosaic supports almost all of the popular Internet protocols, including FTP, NNTP, Gopher, and, of course, HTTP. The Mosaic interface is a window that displays a page of hypertext. The hypertext can combine text, in a variety of fonts, with colour images and even sound and video clips. Hypertext links are displayed as underlined words. By clicking with the mouse on an underlined word, the user causes Mosaic to request a new page of hypertext from a HTTP server. The HTTP server could be anywhere in the world. Thus the user can "surf the net," moving from Canada to Japan to Australia with a click of the mouse button.

The core concept of HTTP is the idea of a URL (Uniform Resource Locator). A URL is a pointer to a file on the Internet. Behind every underlined hypertext word in Mosaic, there is a URL that determines what will happen when the user clicks on the word. A URL is a string of characters, typically beginning with the name of the protocol that is to be used, followed by the Internet address, followed by the directory path, and ending with the file name. (Some of these items are optional and there are others I have left out.) For example, in the URL "http://ai.iit.nrc.ca/cscsi/papers.html", "http" is the protocol, "ai.iit.nrc.ca" is the Internet address, "cscsi" is the directory path, and "papers.html" is the file name.

World Wide Web Sites for AI Workers

If you have Mosaic or another HTTP client, here are two URLs that will be useful starting points for AI workers:

• "http://ai.iit.nrc.ca/ai_point.html": This URL points to a list of AI resources on the Internet. The list includes AI bibliographies, CFPs (Calls For Papers) for AI conferences, AI journals on the Internet, AI research groups with web servers, a subject index for AI resources, AI projects in Canada, and computer science technical reports that are available over the Internet. This is not just a passive list of resources: When the user clicks on the hypertext describing an information resource, the desired information is then delivered to the HTTP client (e.g., Mosaic). The list is maintained by the Knowledge Systems Laboratory of the National Research Council. (I am the current "webmaster" for this web server.)

• "http://ai.iit.nrc.ca/cscsi_point.html": This URL points to information about CSCSI, Canadian AI Magazine, the journal Computational Intelligence, and other AI societies around the world. A growing number of articles from past issues of Canadian AI Magazine are available in hypertext format.

In Mosaic, you can access a URL by pressing the "Open . . ." button at the bottom of the Mosaic window. You then type in your desired URL and press "Open."

Of course, there is a lot more to the web than the above two URLs, but the best way to find out about the web is to get on the web and start surfing. These two URLs will help you to catch the wave.

The World Wide Web as a Research Tool

Beyond net surfing, Mosaic is a very valuable research tool. Permit me to illustrate this point with my own work. I do research in machine learning. My latest project combined a genetic algorithm with a decision tree induction algorithm. With Mosaic, I found and downloaded the source code for an excellent C implementation of a genetic algorithm (Grefenstette's GENESIS). I ordered the decision tree induction algorithm (Quinlan's C4.5) with an order form that I got over the Internet. I wrote some code to join the genetic algorithm with the decision tree induction algorithm, then I tested the resulting software with five medical datasets that I downloaded with Mosaic (from the machine learning dataset repository at the University of California at Irvine). I wrote a paper, summarizing the results of the project, which I submitted to the Journal for Artificial Intelligence Research (JAIR). JAIR is a new electronic journal, distributed over the Internet. I used Mosaic to browse JAIR's instructions for authors. The standard format for a JAIR paper is PostScript, but recently a JAIR paper was written in hypertext (HTML: HyperText Markup Language; the language of the web) and is available from the JAIR web server. If my paper is accepted by JAIR, I will translate it into HTML.

Although JAIR is only a year old, it looks like it is destined to succeed. It has a very impressive editorial board and the typical article is published (electronically, on the Internet) three to four months after it is first submitted. Compare this to the two-year wait that an author can expect with traditional AI journals. Volume One of JAIR contains some very high quality papers, so it is clear that electronic publication does not entail a loss of quality. JAIR papers are available for free — there is no subscription fee for the journal. Using Mosaic, readers can browse the abstracts of papers, then download the PostScript files for articles that seem interesting.

How to Get Mosaic

If you are on the Internet but you do not have Mosaic, here is how you can get it:

```
unix> ftp ftp.ncsa.uiuc.edu
ftp> Name: anonymous
ftp> Password: <your e-mail address>
ftp> cd Mosaic/Mosaic-binaries
ftp> ls
ftp> binary
ftp> get Mosaic-sun.Z
ftp> bye
unix> uncompress Mosaic-sun.Z
unix> Mosaic-sun
```

In the example above, the user has retrieved the Sun version of Mosaic. Mosaic is available for any Unix machine with X Windows. It is also available for PCs with Windows 3.1 and Macs.

There are other HTTP clients, besides Mosaic, that are worthwhile. Mosaic has been the tool of choice for the last few years, but now there is heavy competition from new tools, such as NetScape.

How to Get On the Internet

If you are not on the Internet, you need to contact an Internet access provider. An Internet access provider is a small local company that sells access to the Internet via modem. Check in your yellow pages for an access provider near you.

When you are looking for an Internet access provider, it is important to understand that there are degrees of access. The two important issues are the bandwidth of your connection and the Internet protocols that are available. For bandwidth, you will want at least a 9,000 baud modem. If you wish to run Mosaic, your Internet access provider must support SLIP (Serial Line Interface Protocol) or PPP (Point to Point Protocol). Note that Compuserve, for example, only supports e-mail (SMTP) access to the Internet, although there is a promise of more protocols in Spring of 1995. Delphi supports several protocols, but not SLIP nor PPP. Delphi can give you limited access to the web, using the Lynx HTTP client, which has a text-based interface.

Freenets are sprouting up all over Canada. A Freenet will give you partial Internet access. You will be able to access the web through Lynx (text-only), but not through Mosaic (graphics). The best thing about Freenets is the cost — they're free. Unfortunately, the bandwidth is low, not all protocols are supported, and the dial-in access frequently returns a busy signal.

The cost of Internet access will depend on your local Internet access provider and how much time you spend connected. A typical user might pay about $300 per year. There are rumors that the telephone companies, and also the cable television companies, may become Internet access providers. This competition may make Internet access very inexpensive.

Conclusion

The Internet has been a powerful tool for many years. With Mosaic and similar HTTP clients, the Internet has become a tool that we can all use, without learning obscure command sequences. The Internet is now a vital resource for AI workers . . . and it's fun.

Acknowledgments

Thanks to Joel Martin and Suhayya Abu-Hakima for their comments on an earlier version of this article.

Electronic Publication — Readership Survey

Please reply to Peter Turney's address (listed on Contents page), either by e-mail, fax, or regular mail.

1. Do you have access to the Internet?

   Yes ____ No ____

   If not, do you intend to get access soon?

   Soon ____ Not Soon ____

2. Do you use a World Wide Web client, such as Mosaic, Lynx, or Netscape?

   Mosaic ____ Lynx ____ Netscape ____

   Other ____ None ____

   If not, do you intend to use one soon?

   Soon ____ Not Soon ____

3. Are you in favour of switching from paper publication to electronic publication of Canadian AI Magazine? Some of the benefits would be:
   - decreased or eliminated CSCSI membership fees, due to decreased costs
   - very rapid publication
   - no limitations on the use of colour illustrations (see http://ai.iit.nrc.ca/cscsi/cai/hunter.html or http://www.arc.ab.ca/AITN-ACE.html)
   - no limitations on number of pages

   Some of the negative aspects would be:
   - not all members have Internet access
   - some people prefer glossy magazines to on-line reading or laser printed text

   In Favour ____ Against ____

Comments:

____

____

____

____

____

____

____

____

FEATURE ARTICLES
GROS TITRES

IJCAI: International Jamboree for Canadian Artificial Intelligence

Alan K. Mackworth

The 1995 International Joint Conference on AI, IJCAI '95, will be held in Montréal on August 20-25, 1995. Since IJCAI is being held in Canada for the second time, the editor of this esteemed organ decided, for reasons known only to himself, but with infinite wisdom (editors always have infinite wisdom) to ask me to write a short note on the historical context for IJCAI and Canada. So here it is. We'll start with a little quiz.

IJCAI conferences have been, and will be, held in

- Washington, D.C., USA (1969)
- London, England (1971)
- Stanford, California, USA (1973)
- Tbilisi, Georgia (1975)
- Cambridge, Massachusetts, USA (1977)
- Tokyo, Japan (1979)
- Vancouver, Canada (1981)
- Karlsruhe, Germany (1983)
- Los Angeles, California, USA (1985)
- Milan, Italy (1987)
- Detroit, Michigan, USA (1989)
- Sydney, Australia (1991)
- Chambery, France (1993)
- Montreal, Canada (1995)
- Yokohama, Japan (1997)
- Stockholm, Sweden (1999)
- Somewhere, USA (2001)

Now, as a test for your theory of inductive learning, I would like you, or your program, to induce the simplest rule for planning IJCAI locations based on the historical record. What's the pattern? Don't peek at the answer below.

Here's part of it in Prolog:

```
ijcai_planner(Year, Location) :-
 Year =< 1989,
 Year mod 4 =:= 1,
 continent(Location, north_america).

ijcai_planner(Year, Location) :-
 Year =< 1989,
 Year mod 4 =:= 3,
 continent(Location, Continent),
 Continent =\= north_america.
et cetera
```

You get the idea. It was a trick question: there's no single simple rule. As IJCAI went along it has dynamically invented its own version of the planner. At IJCAI '77, for example, some of us assumed the rule was that the conference alternated between the USA and the rest of the world and made a pitch for IJCAI '79 to be at UBC in Vancouver. At that point the declared rule was clarified to the effect that IJCAI alternated between North America and the rest of the world, so it went to Tokyo. That clarification entitled us to make a successful bid, at IJCAI '79 in Tokyo for IJCAI '81 in Vancouver. Later, as it became clear that IJCAI could better serve its mandate by moving more of its conferences outside North America, the rule was changed so that only every third conference was in North America. And now the lead time for conference planning is six years.

Planning the location involves many factors but the key criteria include local support, both within the country and at the site from the AI community, and the likelihood that the conference will trigger or consolidate growth in the national AI community. Canada will be the first country, besides the USA, to have the honour of hosting IJCAI twice, which is a strong vote of confidence in our AI community.

The Canadian AI community has played an ongoing role in the non-profit society, International Joint Conferences on Artificial Intelligence,Inc.(IJCAII) that co-sponsors the conferences with the local societies (CSCSI and AAAI for IJCAI '95). Don Walker, as the guiding spirit behind IJCAII for many years, successfully built it into a truly international force, while minimizing the bureaucratic inertia that can plague international efforts. IJCAII is run by a minimalist organization consisting of the Board of Trustees, chaired by the current conference chair, and the Secretary-Treasurer. We have declared Ron Brachman, the current Secretary-Treasurer, and Priscilla Rasmussen, who toils behind the scenes, honorary Canadians.

Given the Canadian proclivity for international action it's not too surprising that we have participated in so many ways in IJCAII. Just a few of the connections are mentioned here. The IJCAII Computers and Thought Award to a young AI researcher went to Hector Levesque in 1985. He was the first researcher working outside the US to receive it. In 1987 another Canadian, Johan de Kleer (a UBC and MIT grad working at Xerox PARC) was chosen. In 1993 Ray Reiter was honoured with the IJCAI Award for Research Excellence. People who have served as officers of various IJCAI's include Richard Rosenberg, Local Arrangements Chair (1981), Alan Mackworth, Conference Chair for IJCAI '85, and John Mylopoulos and Ray Reiter, Program Chairs, IJCAI '91. And, of course, for IJCAI '95 the Canadians have taken over: Raymond Perrault is Conference Chair, Renato de Mori is Local Arrangements Chair, Gordon McCalla is Tutorial Chair, and a passel of distinguished Canadians serve on the National Committee.

The impression that Canadians are over-represented both on the Program Committee and in the number of papers presented, compared to the population as a whole and to the size of the community has been confirmed by rigorous analysis; that is, I vaguely recall that someone looked at the statistics once. The idea that there is any causal connection between our presence on the PC and the proportion of papers presented is a scurrilous rumour. The correlation is simply an artefact: both are simple consequences of the excellence of AI in Canada. There's also a rumour that the acronym IJCAI really stands for International Jamboree for Canadian AI. Rumours like these are spread by those paranoids who hold the view that IJCAI has been taken over by 'the Canucks' - a secret society that believes in the power of the Maple Leaf, eh. The paranoids also suggest, based on only a few shreds of evidence such as lapel pins, that there are links between the Canucks and the once-feared, but now toothless, Knights of the Lambda Calculus.

CSCSI is the world's oldest national AI society, founded in 1973. (Both AISB and SIGART are, indeed, older but neither was founded as a national AI society.) So it is only natural that CSCSI should host the world's oldest and still most prestigious AI conference. The choices of Vancouver and Montreal as the two host cities were only natural since they are, as we all know, the hot spots of AI activity in the country.

Over the years many research areas have splintered off from the IJCAI community to form their own specialized workshops and conferences. In 1973, for example, pattern recognition was explicitly excluded causing much sound and fury, and wailing and gnashing of teeth not to mention IJCPR, a pale IJCAI imitation. Theorem proving, representation and reasoning, vision, natural language, constraint programming, robotics, and fuzzy systems have all factored off. Does this mean that IJCAI is obsolete? Absolutely not. Now we realize that carving up the field by function institutionalizes a false view of the design of intelligent agents. IJCAI is, as always, at the centre of the action.

***Alan Mackworth** is a Professor of Computer Science at the University of British Columbia and a Fellow of the Artificial Intelligence and Robotics Program of the Canadian Institute for Advanced Research. He completed a B.A.Sc. in Engineering Science at the University of Toronto in 1966, an A.M. at Harvard University in 1967 and a D.Phil. at Sussex University, England in 1974, when he moved to UBC. He has served as President of the Canadian Society for Computational Studies of Intelligence and as General Chair of the IJCAII Board of Trustees. He has received the Information Technology Association of Canada/NSERC Award for Academic Excellence and the CSCSI/SCEIO Distinguished Service Award. He is currently a Project Leader in IRIS, a Fellow of AAAI, and the Director of the UBC Laboratory for Computational Intelligence. He works on constraints.*

AAAI Calendar Notices

AAAI Spring Symposium Series 1995, sponsored by the American Association for Artificial Intelligence, March 27-29, 1995, Stanford University, California. Contact: Spring Symposium Series 1995, AAAI, 445 Burgess Drive, Menlo Park, CA 94025; 415-328-3123; FAX: 415-321-4457; electronic mail: sss@aaai.org.

IJCAI '95, Fourteenth International Joint Conference on Artificial Intelligence sponsored by the International Joint Conference on Artificial Intelligence Inc.; the American Association for Artificial Intelligence, and the Canadian Society for Computational Studies of Intelligence, August 20-25, 1995, Montreal, Canada. Contact: AAAI, 445 Burgess Drive, Menlo Park, CA 94025; 415-328-3123; FAX: 415-321-4457; electronic mail: ijcai@aaai.org.

IAAI '95, Seventh Annual Conference on Innovative Applications of Artificial Intelligence sponsored by the American Association for Artificial Intelligence, August 20-25, 1995, collocated with IJCAI-95, Montreal, Canada. Contact: IAAI-95, 445 Burgess Drive, Menlo Park, CA 94025; 415-328-3123; FAX: 415-321-4457; electronic mail: iaai@aaai.org.

AAAI Fall Symposium Series 1995, sponsored by the American Association for Artificial Intelligence, November 10-12, 1995, tentatively scheduled to be held at MIT, Massachusetts. Contact: Fall Symposium Series 1995, AAAI, 445 Burgess Drive, Menlo Park, CA 94025; 415-328-3123; FAX: 415-321-4457; electronic mail: fss@aaai.org.

AI Summer in Montreal

C. Raymond Perrault

The Canadian AI community will be host to the world during the Fourteenth International Joint Conference on Artificial Intelligence, which will be held in Montréal, August 20-25, 1995, at le Palais des Congrès. Over 2000 participants are expected.

IJCAI-95 offers Canadian researchers in AI the opportunity to show their international colleagues the exceptional growth and strength of AI in Canada since IJCAI was last held here, in Vancouver in 1981. With the birth of CIAR and its AI and Robotics programme, the funding of the IRIS network, the establishment of PRECARN, and major IJCAI awards to Hector Levesque and to Ray Reiter, no country can boast of a record better than Canada's over the past fourteen years.

IJCAI is jointly sponsored by IJCAI Inc., a non-profit organization, CSCSI, and the American Association for Artificial Intelligence (AAAI). AAAI and a Canadian National Committee organized by CSCSI share the responsibilities for local arrangements. The Canadian Local Arrangements Chair is Renato deMori, of McGill University. Carol Hamilton is the Executive Director of AAAI.

IJCAI will offer a full week of activities, orchestrated by the Program Chair, Chris Mellish, of the University of Edinburgh. The first few days are devoted to tutorials (chaired by Gordon McCalla of the University of Saskatchewan) and workshops (chaired by Tony Cohn of Leeds University), and the rest of the week to the technical program. We expect about 240 papers to be presented, covering the full range of activity in AI. Fourteen invited talks will be given. A video track (chaired by Sam Uthurusamy of General Motors) will allow demonstrations of many systems from research labs around the world. The robot competition/exhibition held at AAAI meetings for the last several years will move its tent to Montreal in 1995. Companies and research labs are invited to show their products and accomplishments at the conference Exhibit Program (Contact ijcai@aaai.org for details). Finally, several activities will be held adjacent to IJCAI. Among them, the Innovative Applications of AI Conference (IAAI) and the Knowledge Discovery in Databases conference.

As Conference Chair for IJCAI - 95, I'd like to invite you all to participate in the preparation of the conference. By the time you read this, it will probably be too late to submit papers (deadline: January 6, 1995), or to propose tutorials and workshops. But there are other ways to help.

IJCAI has managed, longer than almost all other AI conferences, to attract outside sponsorship, to help keep attendance fees down, and to pay for special technical and social events. The organizers of IJCAI-93 in Chambery, France, raised over $400,000 from corporate and government sponsors. Some of this was necessary so that reasonable fees could be charged for a conference in a small town without large convention facilities. And some of it was made possible by early fund raising, at a time when, frankly, AI was more popular with industry than it now appears to be. But it made possible receptions and social activities that have set a new standard for subsequent events, and that Montreal will be hard-pressed to match. Fund-raising activities in Canada have so far been disappointing. If you work for industry, perhaps you could tweak your management to see what they could do for the cause. For further information, you can contact Renato deMori (renato@opus.cs.mcgill.ca).

Conferences are run largely on volunteer labor. The Volunteer Program provides complimentary registration to full-time students, including conference proceedings, in exchange for assisting IJCAI-95 organization in Montreal. For further information, please contact AAAI at volunteer@aaai.org. The deadline for volunteer applications is May 31, 1995. IJCAII will also be coordinating travel awards to assist students, junior faculty, and others who might otherwise not be able to attend. Preference will be given to those with papers accepted in the technical sessions or otherwise directly participating in the conference and who have not previously received an IJCAII Travel Award. For further information, please contact Priscilla Rasmussen (rasmussen@ijcai.org).

For further information about IJCAI-95, the IJCAI organization, and past and future IJCAIs, I invite you to browse the IJCAI WWW site (http://ijcai.org), or to send email to info@ijcai.org, or to me at perrault@ai.sri.com.

The main measure of the success of a conference is the size of its attendance and the richness of the interactions it fosters. The main thing you can do is come! I look forward to seeing you all in Montreal.

***Raymond Perrault** is the General Chair of the 14th International Joint Conferences on Artificial Intelligence, to be held in Montreal August 20-25, 1995. He is also the Chair of the Board of Trustees of IJCAI, Inc (IJCAII), the organizing body of the IJCAI conferences. In his spare time, he is Director of the Artificial Intelligence Center of SRI International, a non-profit research institute in Menlo Park, California.*

Dr. Perrault received a Bachelor's degree in Mathematics from McGill University and a Ph. D. in Computer and Communication Sciences from the University of Michigan. His major research contributions are in natural language processing, speech act theory, discourse, planning, plan recognition, and formal language theory.

Comdale: A Canadian AI Success Story

Aurora C. Diaz

Introduction

Ceci est le récit de Comdale Technologies (Canada) Inc., une compagnie qui a fourni avec succès des logiciels d'IA dans le contrôle des procédés industriels depuis 1986. Comdale est une compagnie canadienne qui développe et met en marché des outils requis pour la surveillance et le contrôle des procédés manufacturiers. La série de produits comprend Comdale/X, Comdale/C, et ProcessVision — chaque produit étant la continuité du précédent. Comdale/X est une coquille de système expert utilisée pour construire des systèmes de consultation hors-ligne. Comdale/C est la continuation de Comdale/X pour les applications critiques en temps-réel. ProcessVision comprend les deux coquilles et les incorpore dans un produit d'acquisition de données et de contrôle intelligent qui fournit différents modules pour l'acquisition et la gestion de donées, la surveillance de procédé, les alarmes, l'interface machine-humain, et le raisonnement expert en temps-réel.

Introduction

This is the story of Comdale Technologies (Canada) Inc., a company that has successfully provided AI software for industrial process control since 1986. Comdale is a Canadian company that develops and markets tools needed to monitor and control both discrete and continuous manufacturing processes. Their suite of products today includes Comdale/X, Comdale/C, and ProcessVision — each product an extension of the previous one. Comdale/X is an expert system shell suitable for building off-line, advisor systems. Comdale/C extends Comdale/X for on-line, time-critical applications. ProcessVision incorporates both shells and extends them into an intelligent SCADA (Supervisory Control and Data Acquisition) package that provides various modules for data acquisition and management, process monitoring, alarming and trending, human-computer interface, and real-time expert reasoning.

History

In 1984, two engineers, a chemical and a mineral processing engineer, decided to build a system that could assist operators in a process plant to efficiently and effectively diagnose process upsets. It was found that much of this help was available in the company operating manuals. However, when faced with an urgent situation, the operators did not have time to browse through the manuals to find the appropriate information. This led the two engineers to build a system that operators could consult during process upsets that would help them to interpret and correct the situation. This system became Comdale/X. Comdale/X, an expert system shell, facilitates the development of off-line advisor systems that are not directly connected to the process environment. A decade ago most AI tools and expert system shells ran on sophisticated workstations and were written in high level languages such as LISP and Prolog. A deliberate decision was made, from the beginning, to develop Comdale/X to run on the standard platforms used in the process plants. For the most part, this meant IBM PC's and compatibles running the DOS operating system. They used the C language to facilitate porting to other platforms. In 1986, Comdale Technologies was formed to market Comdale/X, which today runs on PC's with QNX or OS/2.

Comdale/X

Comdale/X is a rule-based shell that provides both forward and backward chaining and the ability to incorporate user-defined strategies to control inferencing and other aspects of the reasoning process, such as session initiation, session termination, and reporting. The inference control strategies manage the search, focus, and conflict resolution mechanisms. The search mechanism allows for the rule search to be done either depth-first or breadth-first. The focus mechanism, on the other hand, provides various options for finding the optimum solution path, including highest certainty, lowest certainty, highest priority, or lowest priority. When more than one rule may apply, mechanisms including alphabetic, highest priority, and lowest priority may be used to resolve this conflict.

Expert knowledge in the knowledge base is represented as objects, classes, rules, or procedures. Objects are expressed as object-attribute-value triplets and can be of different types, including logical, string, numerical, date, and time. Facets may be added to the triplets to specify information, such as uncertainty, time, fuzzy sets, customized questions and explanations, and sources of information. Classes group similar objects into a hierarchy. Comdale classes have full inheritance capabilities, including multiple inheritance and both public and private attributes. Rules are expressed as if-then-else statements and are used to encode vague concepts, heuristics, mathematical expressions, and time and string expressions. Certainty factors may be attached to rules to reflect the confidence in the conclusion of the rule. Procedures determine the control of rule execution and class/object manipulation.

Comdale/X provides graphical editors and browsers that facilitate the construction of the knowledge base and the debugging of the system. There is a class/object browser that shows the hierarchical relationships between classes and objects and allows for their direct modification. There is also a rule browser that shows the interrelationships among rules and allows rule editing to be done. Debugging facilities include cross-referencing triplets, tracing rules, and watching variables. In addition, support for Microsoft Windows 3.0, QNX Windows, Toolbook, and DV Draw is now provided.

Comdale/C

Late in 1987, in answer to a client's need, there was a move to provide an on-line system that could take data from the process, and use this to automatically provide advice to the operator and, in some instances, to close the loop and send commands to the final control elements in the plant in order to effect control actions. This on-line, real-time version of Comdale/X became known as Comdale/C. The shift from an off-line to an on-line system required the technical decision to migrate from DOS to a real-time operating system. Comdale chose to go with QNX, a multi-tasking, multi-user operating system designed for mission-critical applications. They found that QNX provided the functionality Comdale required to run complex applications on a PC. In addition, QNX is a Canadian product developed by Quantum of Ottawa. This geographic proximity to Comdale made for efficient and productive communication and collaboration between the two companies. Not only did Comdale get good support from Quantum but it also allowed Comdale to provide feed-back to Quantum on the functionality it needed from the operating system and the tools and utilities provided with it.

Comdale/C is designed as an open system composed of individually executable modules that communicate in a client-server or peer-to-peer fashion in a single computer or over a network. Comdale/X forms one of these modules and is used to develop a Comdale/C knowledge base. The inference engine of Comdale/C, known as the controller, subsumes the Comdale/X inference engine. Therefore, an expert system may first be tested using Comdale/X to ensure the correctness of its decisions and then, subsequently, be moved to Comdale/C, where the timeliness of the decisions and the magnitude of the control actions are determined and checked. The controller module augments the Comdale/X inference engine to include mechanisms for reasoning over time. It uses data stored in the Historical and Point databases to reason about past, present, and future events. The controller also has mechanisms that allow the scheduling of future events. It can run in either a cyclic or event-driven fashion or a combination of both. Various external databases, such as Oracle and Informix, can serve as data sources for the controller module but the default data source is the Comdale/C Point Database, which contains up-to-date variable values accessible by any module on the network.

Another module in Comdale/C provides explanations and clarifications needed by the operators and system developers. The ExpertView component of this module discloses the reasoning and decision-making process of the expert system by explaining events, showing knowledge contained in the knowledge base, and expounding operator messages. The ProcessView component, on the other hand, displays process graphics and presents a graphical representation of the analysis and interpretation of process data. Both ExpertView and ProcessView were developed in-house. This decision was influenced by the lack of good, commercial tools in 1987, when the move to on-line systems required such a component for Comdale/C.

ProcessVision

An observation made early on was that there was a lot of interest in using AI to solve problems in industrial process control. However, many of these problems stemmed from the need for a good data acquisition and monitoring package. With this in mind, the company decided to go after the SCADA market with a package that did what other SCADA systems did but that, in addition, integrated other advanced technologies. This package is called ProcessVision, and like Comdale/C, is an open system consisting of independent modules which communicate in a peer-to-peer fashion in a single computer or distributed across a network. Its main modules include the System Modules, the Operator Interface Modules, and the Comdale/C Expert System Extension Modules. The System Modules provide the functionality for system administration, database management, alarming, trending, and batch procedure administration. The Operator Interface Modules include ProcessView and ExpertView together with an editor for ProcessView and a Batch Editor.

Applications of Comdale Products

Comdale products are now used in a range of process plants. They can link with most of the standard hardware found in the plants, including a variety of Programmable Logic Controllers (PLC's), Distributed Control Systems (DCS's), and I/O devices. An example of such a system has been installed at the Dome Mine of Placer Dome Canada Ltd. Dome Mine is a gold mining and mineral processing plant located in Timmins, Ontario. A Comdale system for supervisory control of the grinding circuit was installed in February of 1993. The expert system was developed by Comdale staff and is now being maintained by Dome Mine staff. An interesting comment made by Mr. John Eggert of Dome Mine is that he believes they could "only have designed the knowledge base with difficulty." Even though it looks simple enough to do this given the user-friendly development environment provided with the Comdale products, there is still the need for knowledge engineers to successfully implement a "working" expert system.

There have not been any problems with the system since its installation over a year ago. The only change that has been made was to disable a group of rules that controlled a pump box. Staff at Dome Mine are happy with the Comdale system. Its installation at the plant went very smoothly. Mr. Eggert attributes this to the following factors: (1) the up front work done to optimize and tune the process before the expert system was added, (2) the fact that Dome Mine is highly automated, using the Rosemount System III, a sophisticated DCS, and (3) the high level of expertise of the plant operators. The main benefit derived from the Comdale system is a significant increase in tonnage.

Mr. Eggert credits the addition of the expert system for the increase in tonnage processed by the primary grinding unit, a rod mill, from 161 to 167 dry short tons per hour. As a side effect, they have also learned some interesting things about the process with the use of the Comdale system. For example, a minor property-wide voltage fluctuation caused by the operation of the mine hoist and which, in turn, was causing an unexpected fluctuation in tonnage was brought to their attention after observing the actions of the expert system. In the future, Dome Mine is looking at using Comdale to control the crushing and leaching processes in the plant.

A second example, The Winston Lake Division of Metall Mining Corp., which is a copper-zinc plant located in Schreiber, Ontario installed ProcessVision with Comdale/X and Comdale/C a couple of years ago. Its objective was to provide supervisory control of the grinding and flotation processes in the mineral processing end of the plant. The installation went smoothly enough except for a couple of problems. They had some technical difficulties with the communications driver provided by Comdale that linked the Modicon PLC with the Comdale products, which was resolved. There was also some initial resistance from the operators, which Mr. David Hall attributes to the system having a different interface from what the operators were used to. For example, the operators had problems adapting to the use of the mouse and the mouse-based interface. In spite of overcoming these minor hurdles, this application at Metall has not been very successful.

Although the Comdale products work well, they are not used to their full potential at this installation. The current system does not meet its objective of providing supervisory control of the grinding and flotation processes. It does not provide any advice to the operators on how to better operate the processes nor does it close the loop and actually perform control actions on the final control elements. Instead, it just collects the data from the Modicon PLC and shows this to the operators. Mr. Rich Morrow of Metall says that much of this information is already available without the Comdale system although the operators find the trending information, which is not currently provided by the PLC, quite useful.

Rich Morrow says that they "tried to bite off more than they could chew." They decided to do this installation mostly on their own, including writing their own knowledge base. But, because of other more pressing priorities, they were not able to allocate the necessary manpower to accomplish the task properly. He estimates that at least one person should be dedicated to this to make it work.

Metall is currently working to modify the system to make it more useful to the plant and its personnel. They have called Comdale in to build the knowledge base for them and are currently focusing only on the grinding circuit. Comdale, using its past experience at the plant, some discussions with the plant personnel, and its past experience in writing knowledge bases for grinding circuits, is developing an expert system that will optimize the operation of the rod mill/ball mill grinding circuit. The system should be installed at Metall before the end of the summer of 1994. In addition, in order to speed up the link between the Modicon PLC and the Comdale system, Comdale will be installing a new version of the communications driver.

If things go well with the expert system for the grinding circuit, Metall will consider using the technology for the flotation circuit. Metall's experience brings back the point made by Mr. Eggert. No matter how user-friendly the interface is, building a knowledge base and a working expert system is not easy and requires the expertise of a knowledge engineer.

University Collaboration

Dr. John Meech of the University of British Columbia (UBC) has been a long-time user of Comdale products and his research team, often with support from the National Sciences and Engineering Research Council (NSERC), collaborates closely with Comdale. They get access to the Comdale products in exchange for advice and assistance provided to Comdale in extending these. They use these to teach students about AI and expert systems. Dr. Meech and his team, first at Queen's University and now at UBC, have been involved in the development of over 100 applications using Comdale/X and Comdale/C in various areas, including qualitative mineralogy, mining method selection, flotation reagent advise, fuzzy logic control, froth recognition, steel reheat furnace operations, small hydro plant site selection, waste water process selection, automated speed zone design, and acid rock drainage sampling. More recently, they have built a system for diagnosing mercury pollution in the Amazon. This is currently being field tested in Brazil by graduate students of Dr. Meech.

Meech's group was also involved in developing CRACX, a system for diagnosing cracking problems for continuously cast steel billets built using Comdale products. The development of CRACX was funded in part by a grant from NSERC with matching support from five industrial partners including Alta-Steel, Ivaco, Courtice Steel, Sidbec-Dosco, and Hatch Associates. It is now being used by several steel companies, including Alta-Steel, Manitoba Rolling Mills, Ivaco, and Courtice Steel and is being marketed by Tech-Nexus International Corporation to industry under license from UBC. Like the Dome Mine installation, experts have learned some new things about the process, not only with the use of CRACX, but also during knowledge engineering when transferring their knowledge to a knowledge engineer. The opinion of Dr. Meech is that "Comdale Technologies has played a major role in introducing this technology to the mineral industry and to others."

Reasons for Success

Today there are 400 Comdale systems in use in the market and Comdale has offices in Toronto and Philadelphia and exclusive distributors around the world. The Toronto office employs 25 people. Comdale attributes much of its success to (1) knowing the application domains, (2) the shift from being a provider of expert systems to being a provider of intelligent SCADA systems, and (3) the investment in R&D for continuous development of the products. All three factors boil down to one success criterion — being able to deliver products and services that meet the needs and requirements of the clients.

Comdale began as an AI company founded by non-AI specialists — a mineral processing engineer and a chemical engineer. The company founders were well-aware of the state-of-the-art in their particular domains and also of the outstanding problems here. They learned enough about AI technology to pinpoint the areas where the technology could be useful and to develop this technology into a product. They were able to sell this to the industry because, being members of this community, they could talk the language and had the contacts to help penetrate the industry. This way of working has been carried on in Comdale. Teams that are developing other technologies to be integrated with the product have both technology specialists and application domain specialists. The latter guide the development team in determining what must be achieved and what applications the technology may be suitable for. For example, they are developing neural network technology and the team is composed of both experts in neural networks and experts in industrial process control.

The early applications of Comdale were in mining and mineral processing. Entry here was facilitated by the expertise in the domain of the principals of the company. With a limited amount of marketing and by word of mouth they have been able to make inroads into other industries such as utilities, pulp and paper, cement, and petrochemicals. There was some difficulty in doing this initially when the main offering of the company was an expert system shell. Companies had difficulty fitting in this kind of technology to what they wanted to do. However, with a shift to ProcessVision and SCADA systems and by allying themselves with other companies that sold instrumentation to the industry, Comdale has since gained entry into other industry sectors. SCADA systems are basic tools for any industry sector and this package gave Comdale the ability to enter different markets. ProcessVision is now the base system to which other capabilities may be added, as required. Some additions currently being developed include tools implementing neural networks, genetic algorithms, and machine learning. ProcessVision has shifted the focus of Comdale from being a provider of expert systems to being a provider of advanced technologies integrated and bundled into an intelligent SCADA package.

The main factor that has contributed to the success of Comdale is its ability to deliver products that meet the requirements of its clients. For this reason, much of the company's resources are put into R&D. About 80% of the people work in maintaining and further developing Comdale products. A second group, called Applications Engineering, provides integration, customer technical support, and training. The Technical Services group provides in-house technical support and the rest form part of the Sales, Marketing, and Administration group.

Conclusion

Comdale makes industry and its needs drive the company, including its R&D program. And its track record shows that it has the ability to develop and deliver products required by its clients. That is the main reason why its products are in daily use today and not sitting in a shelf somewhere gathering dust.

Acknowledgments

Thanks to Wayne Thompson of Comdale Technologies (Canada) Inc., John Eggert of Dome Mine, Placer Dome Canada Ltd., John Meech of the Department of Mining and Mineral Process Engineering, The University of British Columbia, and Rich Morrow and David Hall of The Winston Lake Division, Metall Mining for taking time to speak with me about Comdale. This article has been largely based on those interviews together with documentation provided by Comdale and Dr. Meech. Thanks also to Bob Orchard and Peter Turney of the NRC for their comments on the article.

***Aurora Diaz** is a research officer in the Knowledge Systems Laboratory of the National Research Council's Institute for Information Technology. She has been with the NRC since 1988 where her work has focussed on real-time expert systems and their application to the resource industries. More recently she has done work on DAI and agent-based architectures. Aurora has an M.Sc.in Artificial Intelligence from the University of Edinburgh in Scotland and a B.Sc. in Mathematics and Computer Science from Concordia University in Montreal.*

A Personal Report on the Fourth International Conference on Principles of Knowledge Representation and Reasoning

Peter F. Patel-Schneider

The Fourth International Conference on Principles of Knowledge Representation and Reasoning was held 24 - 27 May 1994 amongst the German federal buildings at the Gustav Stresemann Institut in Bad Godesburg, Germany, a suburb of Bonn. There were sessions on Belief Revision, Complexity of Reasoning, Description Logics, Knowledge Sharing and Ontologies, Logics of Knowledge and Belief, Logics of Preference and Utility, Multi-Agent Reasoning, Nonmonotonic Reasoning, Planning, Search, Temporal Reasoning, Theories of Action, and Tractable Reasoning. In total, 55 papers were scheduled for presentation, out of about 270 papers submitted to the conference. Approximately 150 people attended the conference.

The Gustav Stresemann Institut is a dedicated conference facility, containing meeting rooms, bedrooms, a cafeteria, and a Bier pub. This meant that participants did not have to leave the building at all during the conference, except to go to the banquet. I like this arrangement, as it makes it much easier to talk to fellow attendees. The only problem at the Institut was that there were several other groups there also, which reduced the density of KR researchers and also significantly increased the density of tobacco smoke in most natural meeting places.

The vast majority of the papers at the conference were formal, as in past conferences in this series. This was in spite of the call for papers, which stated that the conference was to feature both "the theoretical principles of knowledge representation and reasoning and the relationships between these principles and their embodiments in working systems." The call for papers also encouraged authors to address KR issues involving "real problems" and implemented knowledge representation systems. I think that if the conference is to thrive it must really pick up a significant section of more-practical refereed papers.

One surprise, at least to me, was the largest number of papers on Description Logics. The ten Description Logic papers included formal results on complexity, as in past conferences, but also learning results and uses of Description Logics in Case-Based Reasoning and Plan Recognition.

In contrast, the invited talks and panels were not nearly so formal. Bill Woods gave the keynote address titled "Beyond Ignorance-Based Systems," which I was unable to attend because of thunderstorms in Newark. Jaime Carbonnell gave a talk on Knowledge Representation Issues in Prodigy, detailing the KR issues that arise in this large, integrated AI system. Len Schubert organized a panel on the interaction between Natural Language and Knowledge Representation. (Len had passport problems and unfortunately arrived after his panel.) The most formal invited talk was given by Didier Dubois on the non-standard logics of uncertainty.

Lin Padgham organized a panel on Knowledge Representation and Reasoning Methodologies. (This was one of the best panels I have ever attended.) Most of the participants in the panel, as well as most of the audience, felt that researchers in the area must connect their work to real problems and to work in other areas of AI. A small contingent felt that the field can support long-term theoretical efforts that investigate the properties of KR formalisms without worrying about their proven relationship to real problems. My view is that KR is neither large enough nor sufficiently established to support this sort of research.

The most disturbing facet of the conference was that most of the program committee members did not attend the conference—only 15 of the 52 members were on the registration list. I find this most distressing. If program committee members do not attend the conference, how can they fully check their decisions? How can they expect the conference to remain a good conference? How can they expect others to attend the conference? I foresee the conference on Principles of Knowledge Representation and Reasoning turning into a conference that only presenters attend. What made past conferences great was that I could meet just about everyone in the field I wanted to at the conference.

Personally, I thought that the conference was well worthwhile, but I see a number of disturbing trends that are affecting its viability.

Peter F. Patel Schneider *received his Ph.D. from the University of Toronto in 1987. He has been a member of the AI Principles Research Department at AT&T Bell Laboratories since 1988. From 1983 to 1988 he worked in the AI research group at Fairchild and Schlumberger.*

Peter's research interests center on the properties and use of Description Logics. He has designed and implemented large sections of CLASSIC, a Description Logic-based Knowledge Representation system. He is currently investigating extensions to description logics, including defaults and part-whole relationships — analyzing their properties and utility, and implementing them in CLASSIC.

Machine Learning in Molecular Biology: A review of the Second International Conference on Intelligent Systems for Molecular Biology

Darrell Conklin

The Second International Conference on Intelligent Systems for Molecular Biology (ISMB-94) [1] was held in August on the campus of Stanford University. The conference attracted biologists, chemists, and computer scientists working on a wide variety of AI theories and methods applied to problems in molecular biology, including constraint satisfaction, knowledge representation, and even robotics. However, a significant portion of the papers were concerned with the application and development of new machine learning and discovery techniques for biomolecular data.

The emphasis on machine learning at the conference is not surprising. Protein sequence, structure, and genomic databases are growing at exponential rates and there is a clear need for computational techniques for extracting interesting and useful knowledge from these data. A collection of original machine learning papers drawn from the first conference (ISMB-93) will appear in an upcoming special issue of the *Machine Learning* journal [3]. In his ISMB-94 keynote speech, Bruce Buchanan of the University of Pittsburgh asked "Where is the knowledge we have lost in information?," (from T. S. Eliot's "The Rock") and argued that intelligent systems should be used to handle the overload of information in molecular biology. Machine learning researchers at ISMB-94 were concerned with problems such as learning grammars for RNA and protein structure prediction, discovery of novel protein motifs, and discovery of topological constraints in proteins.

Two tutorials on machine learning preceded the main conference. Pierre Baldi of JPL and Caltech explained the abstract *sequence alignment* task in terms of Bayesian induction and hidden Markov models (HMMs). Sequence alignment equivocates two sequences, potentially allowing for insertions, deletions, and other mutations in the process. An alignment of two or more sequences using an HMM is produced by computing the best path of each sequence through the trained model. David Haussler of UCSC discussed hidden Markov models, stochastic context-free grammars, and graph grammars, emphasizing the EM (Expectation Maximization) algorithm for optimizing grammar production probabilities. Applications to RNA structure prediction were discussed. The power of the basic EM "model-fitting" approach is constrained by the necessarily predetermined architecture of a grammar/model. This shortcoming was addressed by two papers in the main conference proceedings (Fujiwara et al.; Grate et al.).

An emerging theme of the machine learning papers at ISMB-94 was the need for powerful models, or hypothesis spaces, for modelling biomolecular data. Most of these data are *structured*, whether in the form of variable-length sequences or three-dimensional structures. Modelling non-sequential or spatially distant interactions is often critical to success. Thus, attribute-value learning methods have limited applicability to modelling biological phenomena. Some of the more powerful models considered at the conference were HMMs (Fujiwara et al.), stochastic tree grammars (Mamitsuka and Abe), stochastic context-free grammars (Grate et al.), logic programs and other relational models (King et al.; Conklin et al.), and iterative LISP functions (Koza).

A massive amount of data is available to machine learning researchers interested in exploring the fascinating domain of molecular biology. Most of this data is in the form of DNA or protein sequences. A topic of tremendous importance in molecular biology is the discovery, representation, and recognition of *motifs* — short sequence segments — predictive of specific biological functions or three-dimensional structures. Protein motifs can be used for *knowledge-based protein modelling*, where predictions about the structure of an unknown sequence are made by matching against a catalogue of motifs with known structure [2]. While sequence motifs have usually been represented as *consensus patterns* [4], there is a growing interest in probabilistic grammar models. A generalized syntax for these diverse representations of sequence motifs was presented at the conference by Bucher and Bairoch.

Many of the machine learning papers at ISMB-94 were concerned with protein sequence motif discovery. Some of the systems use supervised techniques, where a predetermined classification (i.e., knowledge of function or structure) is supplied with the training sequences. Fujiwara et al. presented an HMM system which learns, in addition to HMM production probabilities, the optimal architecture for the model. The system is applied to the recognition of specific structural regions from sequence data. Grate et al. model the structure of RNA using stochastic context-free grammars. Their system samples a finite set of possible model architectures before invoking an EM training method on a set of RNA sequences with similar structures. The system of Mamitsuka and Abe applies EM to stochastic tree grammars for secondary-structure recognition in protein sequences. Finally, Koza used genetic programming to learn classes of amino acids which are counted in protein segments by iterative LISP functions. These counts are used as parameters to quadratic equations which should be non-zero only for the class of transmembrane protein segments.

When confronted by a mass of uninterpreted sequence data, the biologist often wishes simply to identify the recurrent and statistically significant motifs in the data. Although sequence alignment routines might be classified as "discovery" systems, it is best to reserve that designation for systems which explicitly maintain, evaluate, and filter sets of possible motifs. Three papers at ISMB-94 addressed this unsupervised learning problem. Agarwal and States use existing sequence alignment tools to identify short — possibly mutated — repeated segments, and present an information-theoretic technique for evaluating their significance. Milosavljevic uses dictionary-based data compression techniques to discover and evaluate repeated segments. Finally, Bailey and Elkan use a Bayesian mixture model to discover a single motif within a set of training sequences.

Proteins and RNA molecules do not exist as one- or two-dimensional entities. Rather, they adopt preferred three-dimensional shapes which are largely responsible for their biological function. In addition to sequence data, there exists a large body of three-dimensional structure data. Cohen et al. presented a unified representation and visual query language for three-dimensional motifs. Two papers at ISMB-94 used unsupervised techniques to unearth relationships between the three-dimensional structure and the one-dimensional sequence. Klinger and Brutlag presented a multivariate analysis of the interaction between the side chains of amino acids in alpha-helices. Conklin et al. presented a conceptual clustering technique which discovers interactions between the three-dimensional structure of an amino acid and the three-dimensional structure of an encompassing segment of amino acids.

A few machine learning papers at ISMB-94 were not directly concerned with the representation and discovery of protein motifs. For example, King et al. used the GOLEM inductive logic programming system to discover relational rules for protein beta-sheet topology. These rules predict adjacency, containment, and direction relationships between strands in beta-sheets. The system of Leblanc et al. groups fixed-length DNA sequences using a neural network clustering scheme. Heumann et al. trained a neural network to find complex mappings between fixed-length DNA segments and their ability to act as protein binding sites. Finally, Xu et al. use a neural network to train the weights given to multiple sensors in the GRAIL DNA exon recognition system.

Although ISMB-94 attracted a large and enthusiastic audience of international scientists, there were few researchers based at Canadian universities in attendance. This reflects, in part, the past reluctance of our universities, funding agencies, industries, and investigators to fully support and become involved in collaborative and multi-disciplinary research. The Third International Conference on Intelligent Systems for Molecular Biology (ISMB-95) will be held in July, 1995 at Cambridge, England. The conference promises to be an excellent forum for computer scientists experienced in molecular biology, and for biologists and machine learning researchers wishing to expand their research.

References

[1] R. Altman, D. Brutlag, P. Karp, R. Lathrop, and D. Searls, editors. Proceedings of the Second International Conference on Intelligent Systems for Molecular Biology. AAAI Press, 1994.

[2] T. L. Blundell, B. L. Sibanda, M. J. E. Sternberg, and J. M. Thornton. Knowledge-based prediction of protein structures and the design of novel molecules. Nature, 326:347-352, March 1987.

[3] J. Shavlik, L. Hunter, and D. Searls, editors. Machine Learning Journal Special Issue on Machine Learning in Molecular Biology, 1995. Forthcoming.

[4] J. M. Thornton and S. P. Gardner. Protein motifs and data-base searching. Trends in Biochemical Science, 14:300-304, July 1989.

***Darrell Conklin** is a Ph.D. candidate at Queen's University. He received the B.Sc. and M.Sc. degrees in Computer Science from the University of Calgary. His Ph.D. thesis, entitled "Knowledge Discovery in Molecular Structure Databases," explores the relationships between constitution and conformation in small molecules and macromolecules. His other research interests include knowledge representation, inductive logic programming, computer vision, and algorithmic music modelling and composition. He is an avid musician, and his current love is the violin.*

Become a Member of CSCSI/SCEIO

For more details, see inside back cover.

Canadian Artificial Intelligence

Intelligence Artificielle au Canada

Pure and Applied Research in the Biological Computation Project

Mike Dawson

Sommaire

Le projet BCP est un groupe de laboratoire coordoné par le Dr. Michael Dawson du Département de Psychologie de l'Université d'Alberta. Quoique chapeauté par le département de psychologie, le BCP est interdisciplinaire par nature. Par exemple, des gradués d'Informatique, de Psychologie, et de Philosophie sont actuellement supervisés ou co-supervisés par Dawson et sont tous impliqués en collaboration dans des projets de recherche dans le BCP. Le but premier du BCP est de conduire des recherches pures et appliquées sur les réseaux de neurones artificiels, et d'interrelier ces recherches à des résultats théoriques et empiriques en Science Cognitive.

Introduction

The Biological Computation Project (BCP) is a lab group coordinated by Dr. Michael Dawson of the Department of Psychology at the University of Alberta. Though housed in the psychology department, the BCP is interdisciplinary in nature. For example, graduate students from Computing Science, Psychology, and Philosophy are currently supervised or cosupervised by Dawson, and are all involved in collaborative research projects in the BCP. The primary goal of the BCP is to conduct pure and applied research on artificial neural networks, and to relate this research to empirical and theoretical results in Cognitive Science.

An artificial neural network (ANN) is a computer simulation of a "brain like" system of interconnected processing units. Processing units are typically viewed as being analogous to neurons, and are presumed to operate in parallel. The behaviour of a single processing unit in an ANN can be characterized as follows: first, the unit computes the total signal being sent to it by other processors in the network. Second, the unit applies an activation function to this total signal, in order to adopt a particular level of internal activity. Third, the unit sends a signal to other processors in the network; this signal is a function of the unit's internal activity. The signal that one processor sends to another is transmitted through a weighted connection, which is typically described as being analogous to a synapse.

In general, an ANN can be viewed as a system that generates a desired response to an input stimulus. The pattern of connectivity in an ANN (i.e., the strengths of the connections between various processing units) defines the causal relations between the network's processors, and is therefore analogous to a program in a conventional computer (e.g., Smolensky, 1988). However, in contrast to a conventional computer, the ANN is not given a step by step procedure to perform some desired task. Instead, the network is taught to do the task.

The remainder of this paper provides an overview of the ANN research being conducted by the BCP, and proceeds by describing our work in the following areas: (1) the construction and evaluation of computer models of human motion correspondence, (2) the development of the value unit ANN architecture, (3) the interpretation of network structure, (4) the training of redundant networks, (5) the use of ANNs for medical diagnosis, and (6) the analysis of connectionism's role within cognitive science.

Modeling Motion Correspondence Processing

The roots of the BCP's general interest in artificial neural networks can be found in the study of the motion correspondence problem. This problem is solved whenever a vision system tracks moving objects. Object tracking is problematic because typically the visual information represented in the proximal stimulus (e.g., in retinal stimulation) is impoverished. As a result, there are many different object-tracking interpretations that are consistent with the proximal stimulus. In general, if a system is tracking N different objects, then there are N! different ways of tracking objects, only one of which will be correct (e.g., Ullman, 1979.

Theoretical and experimental research (e.g., Dawson, 1987, 1988, 1989, 1990a; Dawson & Pylyshyn, 1988; Dawson & Wright, 1989) has indicated that solutions to the motion correspondence problem can be achieved by "filtering" proximal stimulus information with three additional constraining properties: the nearest neighbour principle (all things being equal, assign short motion correspondence matches), the relative velocity principle (all things being equal, assign neighbouring elements similar motion correspondence matches), and the element integrity principle (all things being equal, assign matches in such a way that elements do not split apart or fuse together). In general, these results suggest that the visual system should track identities in such a way that (1) the tracking is consistent with the proximal stimulus, and (2) the tracking is also as consistent as possible with the three constraining principles.

Dawson (e.g., 1991) was able to show how these principles could be used as "soft" constraints on the motion correspondence problem by incorporating them into a connectionist network. His model, the "brain state-in-a-sphere," was a variant of Anderson, Silverstein, Ritz, and Jones' (1977) auto associative connectionist network. Processing units in the model represented potential identity matches between elements seen at different positions and at different times. Connection weights were set to represent the three constraining principles. Dawson was able to show that the model generated the same solutions as does the human visual system to a variety of apparent motion displays, which require object identities to be tracked before the illusory motion can be filled in by the visual system. More recent research (Dawson & Wright, 1994; Dawson, Nevin-Meadows, & Wright, 1994) has extended Dawson's original model by building in temporal sensitivity and a preference to assign matches that have the same contrast polarity.

However, Dawson's (1991) motion correspondence model is atypical of most modern ANNs, insofar as the network's connections are "hardwired;" it does not learn. A natural extension of this approach to a problem in computational vision was to consider whether a modifiable network could, through learning, extract different kinds of natural constraints for solving an ill-posed problem. To date, we have not yet addressed this particular question. This is because when our focus was broadened to include trainable artificial neural networks, we made some discoveries that have drawn us away from the original work on motion correspondence. These discoveries are described in the sections that follow.

The Value Unit Architecture

One popular method for training ANNs is the generalized delta rule (Rumelhart, Hinton, & Williams, 1986). With this learning rule, patterns are repeatedly presented to the network, and the network's actual responses are compared to the responses that are desired for these patterns. This comparison involves computing an error term, which can be used to modify the pattern of connectivity in the network in such a way that the network's responses become more and more correct.

The mathematics underlying the generalized delta rule requires that the activation function, used by processing units to compute their internal activity from their net input, be monotonic. This means that when the total signal going into a processor is increased, then the processor's internal activity should never decrease. As a result, most ANNs are based upon a sigmoid-shaped activation function such as the logistic. Adopting Ballard's (1986) terminology, these standard units are called integration devices.

However, there are reasons to believe that neurons would be better modeled with a nonmonotonic activation function (e.g., Ballard, 1986). At the cellular level, many neurons behave as if they are tuned to respond to a narrow range of signals (e.g., a narrow range of light wavelengths, a narrow range of spatial frequencies). If the signal is either lower than or higher than this range, then the cell does not respond. The kind of activation function required to model this type of behaviour is a bell-shaped function, such as a Gaussian.

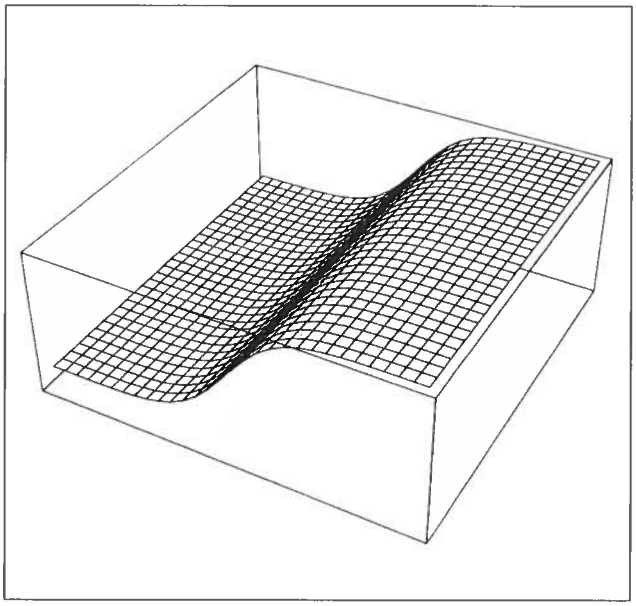

Figure 1. The monotonic, sigmoid shape of the activation function for an integration device. The shadow represents its 2D receptive field.

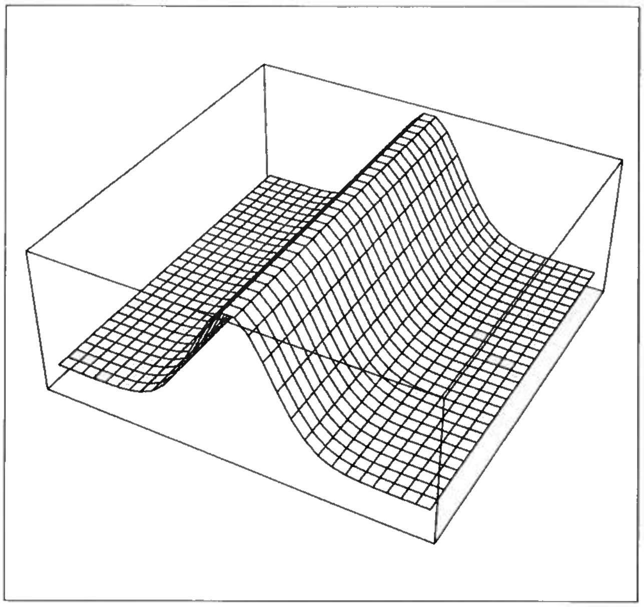

Figure 2. The non monotonic, bell shape of the activation function for a value unit. The shadow represents its 2D receptive field.

Following Ballard's terminology again, a processor that uses such an activation function is called a value unit. Much of the current work at the BCP began with attempts to train networks of value units.

If value units are trained with the standard version of the generalized delta rule, then practical learning does not occur. Instead, the network learns to turn its output units off to every pattern. This is because the mathematics of the generalized delta rule requires monotonic activation functions which are not present in value units. To overcome this problem, an elaborated version of the generalized delta rule was derived (Dawson & Schopflocher, 1992a). For this new learning rule, network error is not only defined in terms of the difference between desired and actual network responses, but also in terms of heuristic information that is used to ensure that the network attempts to turn its output units "on" for some patterns.

The original motivation for developing a learning rule for the value unit architecture was the desire to increase the biological relevance of ANNs (e.g., Dawson, Shamanski, & Medler, 1993). However, after this learning rule was derived, it was soon discovered that value units had many algorithmic advantages over traditional ANNs built with integration devices and trained with the generalized delta rule.

First, value units learn to solve complicated (i.e., linearly nonseparable) problems much faster than do standard networks (Dawson & Schopflocher, 1992a; Dawson, Schopflocher, Kidd, & Shamanski, 1992; Medler & Dawson, 1994; Shamanski, Dawson, & Berkeley, 1994). This is because the new definition of error that was created in the value unit learning rule incorporates a heuristic component that is not used in the standard learning rule. Second, value units appear to better generalize what they have learned on a training set to new patterns, in comparison to standard networks (e.g., Shamanski, Dawson & Berkeley, 1994). Third, value unit networks appear to be better than integration device networks at scaling their performance up to larger versions of the same problem (e.g., Medler & Dawson, 1994). Fourth, the value unit learning rule permits the training of hybrid networks, in which some processors have a nonmonotonic activation function, while other processors have a monotonic activation function (Dawson & Schopflocher, 1992; Dawson, Schopflocher, Kidd, & Shamanski, 1992) The ability to build hybrid networks is attractive, because there is a growing awareness among neuroscientists that brain structure is highly heterogeneous (e.g., Getting, 1989).

Interpreting Network Structure

Mozer and Smolensky (1989, p. 3) have noted that "one thing that connectionist networks have in common with brains is that if you open them up and peer inside, all you can see is a big pile of goo." For researchers interested in using ANNs to model cognitive or perceptual processes, this is unfortunate. For example, McCloskey (1991) makes a strong argument that an ANN cannot be viewed as a theory, a simulation of a theory, nor even a demonstration of a specific theoretical point because of the general inability to interpret the structure of a trained network. Fortunately (and to counter such claims), the value unit architecture has, for binary input patterns, an emergent characteristic that permits a rich interpretation of the internal structure of ANNs.

Consider using a relatively large number of patterns to train an ANN. After training, one could present each pattern once again to the network, and record the activity that each pattern produced in each hidden unit. Then one could use this information to create a jittered density plot for each hidden unit. In such a plot, the horizontal position of each plotted point represents the activation produced by one of the training patterns, and a random vertical jittering is introduced to prevent points from overlapping (Chambers, Cleveland, Kleiner, & Tukey, 1983, pp. 19-21). The purpose of the density plot is to provide some indication of the distribution of activities in the unit that are produced by the training patterns.

Berkeley, Dawson, Medler, and Schopflocher (1994) have found that while the jittered density plots for standard ANN processing units are typically smeared, the same plots for value units are organized into distinct bands or stripes. Furthermore, all the points that fall into a single band share a set of common properties. One can use these properties to identify the specific features in the input patterns that are being detected by the hidden units, and one can also identify

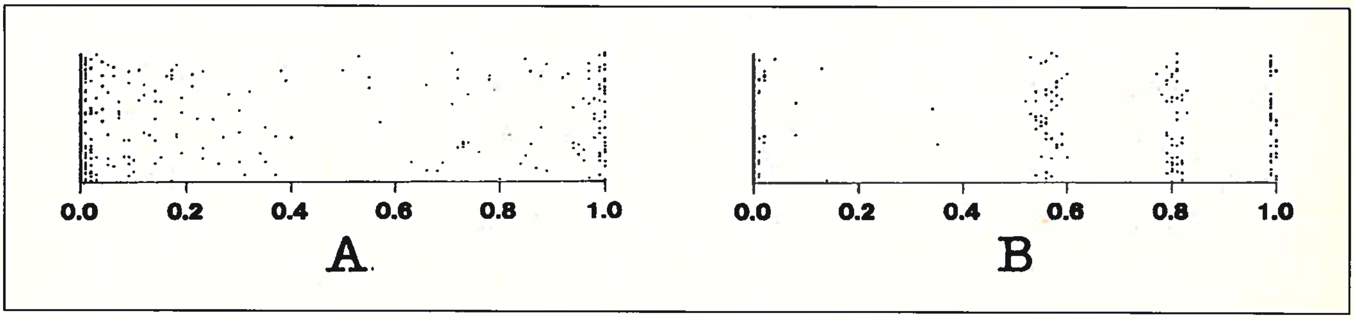

Figure 3. Jittered density plots for hidden units in networks trained on Bechtel and Abrahamsen's (1991) logic problem. The plot on the left is for an integration device; the plot on the right is for a value unit. The latter's distinct bands permit it to be interpreted.

the combinations of these features that are used to mediate the network's output responses.

For example, Berkeley et al. (1994) trained a network on a set of logic problems originally investigated by Bechtel and Abrahamsen (1991). The network is taught to identify the type of logic problem being presented, and to determine whether the problem is valid or invalid. Jittered density plot analysis of hidden value unit activities for this network revealed a tremendous degree of banding; each band reflected a series of important logical properties (e.g., some bands represented the type of conjunction used in the logic problem; others represented balancing between variables in different parts of the problem). They went on to show how these bands could be interpreted to reveal formal rules of logic "in the network's head."

The BCP has also been interested in exploring interpretive techniques to apply to standard ANN architectures. Within perceptual psychology, some researchers adhere to what is known as the neuron doctrine (e.g., Barlow, 1972). According to this doctrine, in order to determine what visual property a neuron is most sensitive to, one must identify the neuron's "trigger feature" — the visual pattern that produces the most activity from the neuron. With the monotonic activation function that characterizes an integration device, it is very easy to identify the trigger feature for a hidden unit simply by inspecting the connection weights that lead into it. Dawson, Kremer, and Gannon (1994) used this technique to demonstrate that under certain conditions, the hidden integration devices in an ANN model of the early visual pathway could evolve biologically relevant receptive fields.

Redundant Networks

The generalized delta rule has proven to be an important and popular procedure for training ANNs. However, it can be characterized as suffering from two major problems. First, it generally is a very slow learning rule. Second, for difficult problems it often produces a local minimum: a configuration of connection weights which cannot be changed by the learning rule, and which result in the network responding incorrectly to some of the input patterns (e.g., Dawson & Schopflocher, 1992a, Table II).

Interestingly, one can take advantage of the local minimum problem in order to speed up learning by exploiting another property of biological networks: redundancy (Medler & Dawson, 1994). Instead of training a single network to solve a problem, a number of similar networks can be trained to solve the same problem, and a decision unit can be trained pool their responses together. The logic of this approach is that if one uses enough variation in initializing the sub networks, then the sub networks will be trained into different local minima (i.e., they will be making mistakes on different input patterns). As a result, the majority of the sub networks will be responding correctly to an input pattern, which permits the decision unit to overrule the incorrectly responding minority of sub networks when responses are pooled.

Initial research on a redundant architecture that uses five sub networks has indicated that redundant networks results in a learning speed-up that justifies the computational expense of training larger networks (Medler & Dawson, 1994). This has been shown for a function approximation problem (training a simulated two-joint robot arm to "reach" to a target) and for a pattern classification problem (training a system to compute whether an input pattern has an even or odd number of "bits" turned on). In both cases, the redundant network learns to solve the problem significantly faster than a single network control, and is more accurate than the control network even when it is given five times the amount of training (to equate the amount of training that has gone into the redundant system). Redundant networks lead to improved performance whether the processors in the units are integration devices or are value units.

Networks For Medical Diagnosis

Medical diagnosis can be viewed as a pattern classification problem; given a set of input measurements, the goal is to categorize a patient as having a particular disorder, or as having no disorder at all. For syndromes (i.e., disorders defined by a constellation of symptoms), this type of pattern classification is almost certainly going to involve linearly nonseparable classes. Furthermore, syndrome diagnosis will likely involve considering nonlinear relationships among multiple variables.

Value unit networks have the potential to be valuable diagnostic tools. First, the research cited above indicates that value unit networks can learn linearly nonseparable classes faster than can standard networks. Second, current results suggest that trained value unit networks will generalize to novel instances better than will standard networks (e.g., Shamanski, Dawson, & Berkeley, 1994), an important consideration if the goal of the system is the accurate diagnosis of new patients.

The potential of value unit networks to diagnosis Alzheimer's disease is currently being investigated (Dawson et al., 1994). The basis for this diagnosis is cerebral blood flow measurements from 14 different brain regions, as determined by single-photon emission computed tomography (SPECT). These measures have been obtained from a large number of patients identified as having probable Alzheimer's disease, as well as from a large number of healthy comparison patients. In one study, a control condition was created by using multiple regression to determine group membership (Alzheimer's vs. healthy comparison) with the 14 SPECT measures used as predictors. The regression equation accounted for only 33.5% of the variance in this diagnostic task. In contrast, networks of value units were able to account for 60% to over 90% of the variance in the data when 10 to 20 hidden units were used in the networks. In fact, a network with 15 hidden units was able to discriminate the two groups perfectly. This performance is extremely encouraging, given the fact that many of the Alzheimer's patients in our sample have a relatively mild manifestation of the disease. Traditional diagnostic methods using SPECT have resulted in accuracy rates as low as 25% for this type of population (Albert & Lafleche, 1991).

Connectionism And Cognitive Science

The preceding sections have indicated that the BCP's research on ANNs has produced some promising results for application: faster learning, easier interpretation, better generalization, and medical diagnosis. However, it is crucial to note that this research was not primarily motivated by the goal to produce such technological advances. Instead, this research has typically been started by asking questions about how the biological relevance of ANNs could be enhanced. Indeed, this has led to design decisions that might seem counterintuitive from an "engineering perspective," but which have led to unanticipated "engineering" advantages (e.g., Dawson & Shamanski, 1994).

The BCP's concern with the relationship between connectionism and cognitive science has been documented in several constructive criticisms (Dawson & Berkeley, 1993; Dawson & Schopflocher, 1992b; Dawson & Shamanski, 1994; Dawson, Shamanski, & Medler, 1993). The general theme underlying these criticisms is that much ANN modeling is driven by technological concerns — building a better mousetrap (i.e., by producing faster algorithms, or architectures that can easily be transferred to silicon chips). There is a growing concern that very little ANN modeling appears to be motivated by theoretical concerns — making models that are strongly related to psychological or physiological phenomena. Most of the BCP's simulation research has been motivated by these latter concerns, and has taken some small steps to enhance the role of connectionism within cognitive science.

References

Albert, M.S., & Lafleche, G. (1991). Neuroimaging in Alzheimer's disease. The Psychiatric Clinics of North America, 14, 443-459.

Anderson, J.A., Silverstein, J.W., Ritz, S.A., & Jones, R.S. (1977). Distinctive features, categorical perception, and probability learning: Some applications of a neural model. Psychological Review, 84, 413-451

Ballard, D. (1986). Cortical structures and parallel processing: Structure and function. Behavioural and Brain Sciences, 9, 67-120.

Barlow, H.B. (1972). Single units and sensation: A neuron doctrine for perceptual psychology. Perception, 1, 371-394.

Bechtel, W., & Abrahamsen, A. (1991). Connectionism and the mind. Cambridge, MA: Basil Blackwell.

Berkeley, I.S.N., Dawson, M.R.W., Medler, D.A., & Schopflocher, D.P. (1994). Density plots of hidden value units reveal interpretable bands. Neural Computation, under editorial review.

Chambers, J.M., Cleveland, W.S., Kleiner, B., & Tukey, P.A. (1983). Graphical methods for data analysis. Belmont, CA: Wadsworth International Group.

Dawson, M.R.W. (1987). Moving contexts do affect the perceived direction of apparent motion in motion competition displays. Vision Research, 27, 799-809.

Dawson, M.R.W. (1988). The cooperative application of multiple natural constraints to the motion correspondence problem. In R. Goebel (Ed.) Proceedings of the Seventh Canadian Conference On Artificial Intelligence. Edmonton, AB: University of Alberta Printing Services.

Dawson, M.R.W. (1989). Constraining tag-assignment from above and below. Behavioral and Brain Sciences, 12, 400-402.

Dawson, M.R.W. (1990a). Apparent motion and element connectedness. Spatial Vision, 4, 241-251.

Dawson, M.R.W. (1990b). Training networks of value units: Learning in PDP systems with nonmonotonic activation functions. Canadian Psychology, 31(4), 391. (Invited abstract.)

Dawson, M.R.W. (1991). The how and why of what went where in apparent motion: Modeling solutions to the motion correspondence problem. Psychological Review, 98, 569-603.

Dawson, M.R.W., & Berkeley, I. (1993). Making a middling mousetrap. Behavioural and Brain Sciences,16, 454-455.

Dawson, M.R.W., Dobbs, A., Hooper, H.R., McEwan, A.J.B., Triscott, J., & Cooney, J. (1994). Artificial neural networks that use SPECT to identify patients with probably Alzheimer's disease. European Journal of Nuclear Medicine, in press.

Dawson, M.R.W., Kremer, S., & Gannon, T. (1994). Identifying the trigger features for PDP models of the early visual pathway. In R. Elio (Ed.) Proceedings of the Tenth Canadian Conference On Artificial Intelligence. Palo Alto, CA: Morgan Kauffman (pp. 115-119).

Dawson, M.R.W., Nevin-Meadows, N., & Wright, R.D. (1994). Polarity matching in the Ternus configuration. Vision Research, in press.

Dawson, M.R.W., & Pylyshyn, Z. (1988). Natural constraints on apparent motion. In Z.W. Pylyshyn (Ed.) Computational processes in human vision. Norwood, NJ: Ablex.

Dawson, M.R.W., & Wright, R.D. (1989). The consistency of element motion affects the visibility but not the direction of apparent movement. Spatial Vision, 4, 17-29.

Dawson, M.R.W., & Wright, R.D. (1994). Simultaneity in the Ternus configuration: psychophysical data and a computer model. Vision Research, 34, 397-407.

Dawson, M.R.W. & Schopflocher, D.P. (1992b). Autonomous processing in PDP networks. Philosophical Psychology, 5, 199-219.

Dawson, M.R.W. & Schopflocher, D.P. (1992a). Modifying the generalized delta rule to train networks of nonmonotonic processors for pattern classification. Connection Science, 4, 19-31.

Dawson, M.R.W., Schopflocher, D.P., Kidd, J., & Shamanski, K.S. (1992). Training networks of value units. Proceedings of the Ninth Canadian Conference on Artificial Intelligence. (pp. 244-250).

Dawson, M.R.W., & Shamanski, K.S. (1994). Connectionism, confusion and cognitive science. Journal of Intelligent Systems, in press.

Dawson, M.R.W., Shamanski, K.S., & Medler, D.A. (1993). From connectionism to cognitive science. In L. Goldfarb (Ed.) Proceedings of the Fifth University of New Brunswick Symposium On Artificial Intelligence, Fredericton, NB: UNB press (pp. 295-305).

Getting, P.A. (1989). Emerging principles governing the operation of neural networks. Annual Review of Neuroscience, 12, 185-204.

McCloskey, M. (1991). Networks and theories: The place of connectionism in cognitive science. Psychological Science, 2, 387-395.

Medler, D.A., & Dawson, M.R.W. (1994). Training redundant artificial neural networks: Imposing biology on technology. Psychological Research, in press.

Mozer, M.C., & Smolensky, P. (1989). Using relevance to reduce network size automatically. Connection Science, 1, 3-16.

Rumelhart, D.E., Hinton, G.E., & Williams, R.J. (1986). Learning representations by backpropagating errors. Nature, 323, 533-536.

Shamanski, K.S., Dawson, M.R.W., & Berkeley, I.S.N. (1994). The effect of linear separability on learning speed and generalization in monotonic and nonmonotonic PDP networks. Journal of Artificial Neural Systems, under editorial review.

Smolensky, P. (1989). On the proper treatment of connectionism. Behavioural and Brain Sciences, 11, 1-74.

Ullman, S. (1979). The interpretation of visual motion. Cambridge, MA: MIT Press.

Acknowledgments

This paper was supported in part by Research Grant A2038 from the Natural Sciences and Engineering Research Council of Canada. Thanks to Nancy Digdon for her comments on the manuscript. Address correspondence to Dr. Michael Dawson, Biological Computation Project, Department of Psychology, University of Alberta, Edmonton AB CANADA T6G 2E9. E-mail: mike@psych.ualberta.ca

---

Michael R.W. Dawson *is an associate professor of psychology at the University of Alberta. He is also an adjunct member of the departments of philosophy and computing science. He received his Ph.D. in psychology from the University of Western Ontario in 1986. He was a member of the psychology department at York University in 1986-87 and arrived at the University of Alberta in 1987. He has published over 33 journal articles, 11 abstracts, and 6 technical reports, and is currently writing a book on the foundations of classical cognitive science.*

Call for Papers: "Machine Learning"

The Autumn 1995 issue of ***Canadian Artificial Intelligence*** magazine will be a Special Theme issue devoted to Machine Learning. If you are interested in contributing an paper on this topic, please note the deadline for receipt of papers is July 15, 1995.

If you intend to submit a paper, please inform the editors in advance, so that space can be reserved for your paper. Papers that are not on the above items are also welcome, as are news items, or conference reviews

Contributions for the Spring/Summer 1995 issue of Canadian Artificial Intelligence magazine must be received by **March 30, 1995.** Contributions for the Autumn 1995 issue are due *July 15, 1995.*

Please send your contribution, electronic preferred, with an abstract, photograph, and short biography to:

Dr. Peter Turney, Co-Editor,
Canadian Artificial Intelligence Magazine
Knowledge Systems Laboratory
Institute for Information Technology
National Research Council Canada
M-50 Montreal Road
Ottawa, Ontario, Canada
K1A 0R6 or — peter@ai.iit.nrc.ca

Canadian Artificial Intelligence
Intelligence Artificielle au Canada

Jean-Claude Gavrel

PRECARN Launches Its New $50 Million Research Program

On January 6, 1995, Dr. Jon Gerrard, Secretary of State for Science, Research, and Development, confirmed Industry Canada's contribution of $19.4 million over the next five years, towards the second phase of research initiatives led by PRECARN.

"PRECARN's success in building partnerships between industry, universities, and governments to develop and disseminate advanced technology exemplifies the collaborative approach that Canada must pursue to meet the challenges of the global economy," said Dr. Gerrard. "PRECARN is leading the development of a world-class intelligent systems research capability in Canada that will translate into increased competitiveness for businesses of all sizes in a wide-range of industry sectors."

PRECARN's research program will continue to feature long-term industry-based collaborative research projects, but will also highlight closer-to-market research initiatives, with a strong emphasis on technology diffusion. The Request for Proposals (RFP) for this second wave of projects is expected to be issued in late February 1995. With a deadline for submissions of early May, it is anticipated that the successful proposals will be announced at the Fifth Annual IRISPRECARN Conference, June 13-15, 1995, in Vancouver, British Columbia.

"With Industry Canada's commitment now confirmed, PRECARN is gearing up for a $50 million research program. Our challenge is to build on the success of phase 1 - success in creating new partnerships and working relationships among governments, industry, and universities, and success in developing new intelligent systems technologies," says PRECARN's Interim President Dr. J. Ron McCullough. "Our aim over the next five years is to make sure that these very specific, concrete benefits are offered to all our partners — in industry, universities and governments."

The second wave of PRECARN's research program is designed to respond to a range of industrial needs. It will include a mix of:

Long-term projects, to address complex industrial problems where some element of new basic research is required. These projects are high risk and would typically involve three or more Members and one or more research institutions. The output of such projects will normally be new solutions in the form of research results.

Medium-term projects, to address more focussed industrial problems and where some initial research activity is already underway. These projects are medium risk and would involve at least two Members and a research group already working on the research problems. The output of these projects would be new technologies having been tested in at least one industrial setting.

Short-term projects, where industry wishes to take the results of a basic research activity and further develop and test it in an industrial environment. These projects are medium risk and would normally involve at least two Members and participants from the original basic research group. Output from these projects would be proof-of-concept prototypes.

"Phase 2 of PRECARN's research program will place a greater emphasis on attracting Small to Medium-Sized Enterprises (SMEs). The importance of the role that SMEs play in the Canadian economy has been recognized by governments and industry, and it is crucial that their expertise be fully utilized to place Canada on an international scale," says Dr. McCullough.

PRECARN will also continue to build on the research accomplishments of the Network of Centres of Excellence that it manages, the Institute for Robotics and Intelligent Systems (IRIS). With the expertise of over 300 university researchers and their graduate students in 24 Canadian universities, this $50 million, eight-year research program, has several projects ready to be transferred into industry-based research initiatives.

To obtain a copy of the Request for Proposals, or for more information on PRECARN, please contact Lise McCourt, Tel: (613) 727-9576, Fax: (613) 727-5672, E-Mail: mccourt@precarn.ca

Canadian Artificial Intelligence
Intelligence Artificielle au Canada

BOOK REVIEWS

Representing uncertain knowledge: An artificial intelligence approach Paul Krause and Dominic Clark (Imperial Cancer Research Fund, London) Oxford: Intellect Books and Dordrecht: Kluwer Academic Publishers, 1993, ix+277pp; hardbound, ISBN 0-7923-2433-1, no price listed; paperbound, ISBN 1-871516-17-X, 14.95 pounds

Reviewed by
Eric Neufeld
University of Saskatchewan

Representing Uncertain Knowledge is a clear and compact overview of some prevailing schools of thought within the uncertainty subcommunity. It requires minimal mathematical background and would be useful for getting up to speed in the area. It would also make a good supplementary text in a course on knowledge representation, although there are no exercises.

Uncertainty itself is a vague term; it isn't hard to cast many or even most AI problems as reasoning under uncertainty. What I think of as the uncertainty community developed largely apart from the mainstream AI community and formed a nucleus in the early 1980s around the Uncertainty in Artificial Intelligence Workshops, which brought together those studying Bayes nets, certainty factors, the Dempster-Shafer theory of evidence, and fuzzy logic. From this meeting emerged a lively polemic on the foundations of uncertain inference and a re-examination of the foundations of classical probability theory in particular. Members of this community challenged the logicist tradition in the mainstream, and later went on to harmonize the logic of statistical inference with nonmonotonic reasoning. Most of the topics discussed in *Representing Uncertain Knowledge* come from here.

Although uncertainty formalisms evolved historically from efficient but crude heuristics to current sophisticated probabilistic techniques, this book takes an opposite tack, presenting first probability as a tried-and-true standard against which to judge subsequently the heuristic formalisms. After discussing different kinds of uncertainty (probability, imprecision, vagueness, incompleteness, error, and lack of confidence), the authors briefly introduce the foundations of probability à la Cox followed by a discussion of the "myth of modularity in expert systems," that is, the notion that an expert system might consist of completely independently firing rules. This hopeful thought is dismissed by pointing out that probability statements can't be composed. The extreme opposite case is when variables are completely dependent and relations can only be expressed by the equivalent of a joint probability distribution, requiring exponential space and time to represent and evaluate.

The practical solution between the extremes assumes reasonable independence relations among variables allowing the joint to be factored into smaller local distributions that make computational tractable. Bayes nets are introduced as probability with an effectively computable face. For brevity, only the algorithms of Lauritzen and Spiegelhalter are given.

The next three chapters present the salient features of certainty factors, the Dempster-Shafer theory of evidence, and possibilistic and fuzzy logic, respectively. The treatment is fair and even, yet not dry. On one hand, non-probabilist formalisms are not subjected to polemics. The authors recall the statement of Pearl and Shafer that not all problems in AI lend themselves to probability. On the other hand, the authors clearly identify well-known trouble spots, such as Zadeh's critique of Dempster's combination rule in the presence of conflicting arguments.

The last third of the book discusses symbolic approaches to uncertain reasoning. This is a good idea, but it came as a surprise that there is no treatment of qualitative probability or the logic of statistical inference, yet a whole chapter is devoted to nonmonotonic logic. However, this probably reflects the composition of the European Uncertainty Conference.

Some similar quibbles aside, *Representing Uncertain Knowledge* is a readable fast track into uncertain reasoning that gives the reader the basic information as well as a sense of some of the lively debates in the area.

***Eric Neufeld** is an Associate Professor in the Department of Computational Science at the University of Saskatchewan. He is interested in probabilistic and statistical approaches to AI problems that have traditionally been handled by logics, grammars, or heuristics.*

Hybrid neural network and expert systems *Larry R. Medsker* (The American University) Boston: Kluwer Academic Publishers, 1994, x+240 pp; hardbound, ISBN 0-7923-9423-2, US$87.50

Reviewed by
Lawrence A. Bookman
Business Extracts International

Given the recent strong interest in both neural networks and their entrance into the commercial world, and the ongoing interest in expert systems, it seems only natural to combine both of these technologies in the hope of getting the best of both. Larry Medsker has written a practical book that successfully gives the flavor of this combined technology. The book is subdivided into three parts. Part I provides an introduction to the field and discusses the fundamentals of hybrid systems. Part II presents five case studies that combine neural networks and expert systems. Part III describes general guidelines for developing hybrid systems and some existing tools and development systems. It concludes with a discussion of the future of this combined technology.

The case studies described in Part II represent the strongest part of the book. They are a collection of research projects done in mostly commercial research and development environments. The scope of these projects range from hybrid systems for window glazing design, chemical control, nuclear plant monitoring, and image interpretation, to multiple target recognition. What is nice about these studies is that each one is well described, attacks a real problem, presents the strengths and weaknesses of both the neural network and expert system approaches, and evaluates their proposed solution.

Part I, although not as strong as Part II, presents a simple, quick introduction to neural and symbolic systems along with some useful pointers into the research community. However, this section does contain some minor flaws. Specifically, the first two chapters of this section are somewhat sketchy and lacking in specifics. Personally, I would like to have seen a little more synthesis and analysis of the aforementioned research pointers, but maybe that is beyond the scope of the book. The last chapter in this section presents a nice description of several possible integration strategies. Part III includes a good general how-to guide for integrating neural networks and expert systems from the concept phase to application design. A stronger presentation of this material could have been achieved if more specifics were provided to tie this material to Part II.

The book could have used some copyediting, as it contains misspellings, pointers to nonexistent figures and to incorrect figures, in addition to some sloppy writing. Overall, I would recommend this book mostly because of the strength of its excellent case study descriptions.

__Lawrence A. Bookman__ is President and co-founder of Business Extracts International. His books Trajectories through knowledge space: A dynamic framework for machine comprehension and Computational architectures integrating neural and symbolic processes: A perspective on the state of the art *(co-edited with Ron Sun) were published by Kluwer Academic Publishers in 1994.*

Decision-analytic intelligent systems: Automated explanation and knowledge acquisition *David A. Klein* (Boston University) Hillsdale, NJ: Lawrence Erlbaum Associates, 1994, xi+220 pp; hardbound, ISBN 0-8058-1105-2, US$49.95

Reviewed by
Darren Meister
Carleton University
and
Stan Szpakowicz
University of Ottawa

The book presents a methodology, called Interpretive Value Analysis (IVA), "for explaining and refining decision-analytic choices automatically." The goal of IVA is to increase the transparency of a class of decision-analytic solutions, namely those using multiattribute value theory in a formal and structured manner. IVA comprises an interpretation formalism for presenting and discussing value-based choices, and strategies for explaining and refining value-based choices. The decision framework, based on the work of Keeney and others, uses numerical preferences and numerical criterion importance.

The author starts out by clearly stating his goals and introducing the reader to the area. He gives the motivation for the development and implementation of IVA, and stresses the importance of formalization and transparency in systems based on decision analysis. VIRTUS, a system based on IVA, illustrates the use of the framework. Next, the components of IVA are presented through example and thorough formal definition. Three example applications run throughout the text. The book concludes with comments on the strong points of IVA and its research and decision-making future.

The author has nicely assembled several ideas and techniques from decision analysis and AI, but the book offers little by way of advancement in either field. In particular, important problems in decision analysis, such as preference elicitation or the validity of criterion weights, are only briefly touched upon and no solutions are suggested.

While the aim of the book is laudable and the actual work quite solid, it is difficult to extract knowledge from it. It feels too much like a thesis. It is hard to read — over-referenced, over-quoted, hugely over-footnoted. There is no doubt that it was a worthy dissertation, but even a very respectable dissertation need not make a good book. The tenets are simple and not unappealing, and the skill with which they have been worked out is considerable, but the packaging is verbose. Almost everything in the central chapters is spelled out *in toto*, just to show that the Author can do it; very little is left to the reader's imagination. Chapter 9 positively screams "This is a student work!".

This would not be a problem except for the fact that the focus is narrow and the research results are not mature. Sufficient links to the research world around IVA are not drawn; this may leave an interested scientist thinking, "A good research goal, but how and why would I want to bother using this?" One problem is that the book deals with only one method of decision analysis, without addressing other methods that aim to improve the transparency of the process in a formal setting (for example, methods based on non-compensatory preference structures and even the analytic hierarchy process). In his notes on future research, the author sets the goal of extending IVA to interpretive decision analysis, but there are no hints how this might be accomplished, and no references (a rare oversight) to what comprises decision analysis in the author's opinion.

The examples have been dressed up for presentation, and for that matter dressed up as three different applications. They are, however, quite similar in that they employ the same vector-shuffling machinery; the good-looking phrases that conceal numbers are irrelevant. This does not help to show how IVA can be used in diverse settings.

The empirical results reported in the book (Section 4.1) consist of fragments of interviews with specialists and laymen. They have not been ordered in any manner other than being squeezed into four visibly preconceived categories. It is very neat, and the presence of a single fundamental concept of compellingness is elegant, though this would hardly be useful for an average reader and user. The presentation is, however, too explicit. Some completions of lists of measures, functions, and predicates are quite obvious (for example, Section 4.3.5) and could have been left unspelled — the reader should be trusted with sufficient acuity. The book's intent is to be thought-provoking and informative, but it is not easy when the reader's mind has been numbed with details. On the other hand, a list of names may have been added; we had trouble locating in Chapter 4 names mentioned in Chapter 5. By the way, the predicate INCONSEQUENTIAL? (Section 5.2.1) is undefined, unless it harks back to one value in an arbitrary list of values for compellingness.

While the value of formalization is stressed, at times this goal is lost. For one, the definition of formal systems is not formal. The discussion of pruning a tree of values for a more succinct presentation to the user (Section 5.2.1) is interesting, but we stumble at the very beginning; in vertical pruning, the crucial concept of interest is "a matter of taste." So much for formality? Examples of horizontal pruning are nice, though they come out of the blue.

There follows a very interesting discussion of difference-function traversal as a mechanism to generate explanations; it is based on a model that has been skillfully adapted for this work. Though the example on pages 127-129 is perhaps too long, this is the best part of the book. It accounts for part of the book's title whose other parts are not entirely justified by the contents ("intelligent systems" and "knowledge acquisition" come to mind).

The discussion of interactive diagnosis and repair (Section 6.2.1) is interesting. On the other hand, though the example illustrating acquisition (Section 6.3) is very long, it does not begin to tell us how the knowledge engineer is expected to "impromptu" the missing factors and their complete numerical profile.

Chapter 7 may apparently be skipped by the reader who lacks background in computing. It is expendable even for such a reader; VIRTUS's reliance on the KEE system makes it unattractive for a vast majority of potential users.

Section 8.1 contrasts the author's approach with that found in other, mostly much earlier, systems: among these, MYCIN is used as backdrop. This section, while indispensable in a thesis, is of little value in a book. The claims regarding transparency are not defensible at all, insofar as no definition of transparency has been given. Transparency is described as "an intuitive framework for interpreting results and for systematically modifying its parameters." This must be true for the user of the system, not just the designer. It is premature to say that IVA is a "more" transparent system than MYCIN without empirical evaluation using domain decision-makers. For example, there is no evaluation of whether the preference elicitation method introduces bias, an important issue in decision making.

What is the author proposing, really? Not a tool that could be widely used in practice (though it may have the makings thereof). Not a simple, catching framework — this is too baroque. Maybe a system of functions and predicates dealing with values, although few of them can be simultaneously employed in one system, as the example of VIRTUS shows clearly.

Part of the problem with the book is that the title leads you to believe that it will be a discussion of decision-analytic intelligent systems *in general*, whereas the content and style try to convince you of the superiority of IVA and VIRTUS. This is a disappointment. Is it worth reading? For some — but it is not an easy read. Is it worth having on the shelf? Yes, for the references, and perhaps yes, as a reference to respectable scholarly work. As a thought-provoking study — no. As a supplementary reading in a course — not at all. Notwithstanding all that, a different book (part pruned, part added to) would have made an interesting reading.

Darren Meister *is an NSERC Post-Doctoral Fellow currently at the School of Business, Carleton University. He received his PhD in Management Sciences from the University of Waterloo in 1993. His research interests are multicriteria decision making, integration of AI and decision-making models, and managerial issues involving computer-based support.*

Stan Szpakowicz *is a professor in the Department of Computer Science, University of Ottawa. He did his PhD in computational linguistics in Poland in 1978. His current interests include AI (natural language processing, knowledge representation) and decision modelling, analysis, and support. He has published four books and over seventy research papers.*

BOOKS RECEIVED

Books marked with a + in the list below are scheduled for review in a future issue. Reviewers are still sought for books marked with a *. Readers who wish to review books for *Canadian Artificial Intelligence* should write, outlining their qualifications, to the book review editor, Graeme Hirst, Department of Computer Science, University of Toronto, Toronto, Canada M5S 1A4, or send electronic mail to gh@cs.toronto.edu or gh@cs.utoronto.ca. 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 Artificial Intelligence* 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.

Parallel natural language processing *Geert Adriaens and Udo Hahn (editors)* (University of Leuren and University of Freiburg) Norwood, NJ: Ablex Publishing Company, 1994, vi+467 pp; hardbound, ISBN 0-89391-869-5, US $79.50 (US $37.50 prepaid)

The language of first-order logic (third edition, revised and expanded), including the Macintosh version of Tarski's World 4.0 *Jon Barwise and John Etchemendy* (Indiana University and Stanford University) Stanford: Center for the Study of Language and Information (CSLI lecture notes 23), 1993, xiv+313 pp and 3.5 inch diskette; paperbound, ISBN 0-937073-99-7, no price listed

+Computational organization theory *Kathleen M. Carley and Michael J. Prietula (editors)* (Carnegie Mellon University) Hillsdale, NJ: Lawrence Erlbaum Associates, 1994, xvii+318 pp; hardbound, ISBN 0-8058-1406-X, US $89.95

Music processing *Goffredo Haus (editor)* (University of Milan) Madison, WI: A-R Editions Inc. (The computer music and digital audio series, edited by John Strawn, volume 9), 1993, x+403pp; paperbound, ISBN 0-89579-268-0, US $39.95

***The Gödel programming language** *Patricia Hill and John Lloyd* (University of Leeds and University of Bristol) Cambridge, MA: The MIT Press (Logic programming series, edited by Ehud Shapiro), 1994, xx+348 pp; hardbound, ISBN 0-262-08229-2, US $45.00

***Prime Time freeware for AI, Issue 1-1** *Mark Kantrowitz (editor)* Sunnyvale, CA: Prime Time Freeware, 1994, 220 pp and two CD-ROMs; ISBN 1-881957-11-X, US $60.00

Verbmobil: A translation system for face-to-face dialog *Martin Kay, Jean Mark Gawron, and Peter Norvig* (Xerox Palo Alto Research Center, SRI International, and Sun Microsystems Laboratories) Stanford: Center for the Study of Language and Information (CSLI lecture notes 33), 1994, viii+235 pp; hardbound, ISBN 0-937073-96-2, no price listed; paperbound, ISBN 0-937073-95-4, no price listed

***Advances in genetic programming** *Kenneth E. Kinnear, Jr (editor)* (Adaptive Computing Technology) Cambridge, MA: The MIT Press, (Complex adaptive systems series), 1994, xiii+519 pp; hardbound, ISBN 0-262-11188-8, US$45.00

Image and brain: The resolution of the imagery debate *Stephen M. Kosslyn* (Harvard University) Cambridge, MA: The MIT Press, 1994, viii+516 pp; hardbound, ISBN 0-262-11184-5, US $45.00

***Genetic programming II: Automatic discovery of reusable programs** *John R. Koza* (Stanford University) Cambridge, MA: The MIT Press, (Complex adaptive systems series), 1994, xx+746 pp; hardbound, ISBN 0-262-11189-6, US $45.00

Computer program construction *Ali Mili, Jules Desharnais and Fatma Mili* (University of Ottawa, Université Laval, and Oakland University) New York: Oxford University Press, 1994, xix+379 pp; hardbound, ISBN 0-19-509236-8, Cdn $81.50

Dictionnaire encyclopédique de pragmatique *Jacques Moeschler and Anne Reboul* (Université de Genève and Centre de Recherche en Informatique de Nancy) Editions du Seuil, 1994, 579 pp; hardbound, ISBN 2-02-013042-4, FF 380

Participating in explanatory dialogues: Interpreting and responding to questions in context *Johanna D. Moore* (University of Pittsburgh) Cambridge, MA: The MIT Press (The ACL - MIT Press series in natural language processing, edited by Aravind K. Joshi, Karen Sparck Jones, and Mark Y. Liberman), 1995, xiv+352 pp; hardbound, ISBN 0-262-13301-6, US $49.95

***Circuit complexity and neural networks** Ian Parberry (University of North Texas) Canbridge, MA: The MIT Press (Foundations of computing series), 1994, xxxi+270 pp; hardbound, ISBN 0-262-16148-6, US $35.00

Sémantique pour l'analyse de la linguistique à l'informatique *François Rastier, Marc Cavazza, and Anne Abeillé* (CNRS, Thomson-CSF/LCR, and Université de Paris VII) Paris: Masson (Collection sciences cognitives), 1994, xii+240 pp; broché, ISBN 2-225-84537-9, FF 180.00

Language computations *Eric Sven Ristad (editor)* (Princeton University) [DIMACS Workshop on Human Language, 20 - 22 March 1992] Providence, RI: American Mathematical Society (DIMACS series in discrete mathematics and theoretical computer science, volume 17), 1994, xiv+198 pp; hardbound, ISBN 0-8218-6608-7, US $60.00

Compositional translation *M.T. Rosetta (pseud.)* Dordrecht: Kluwer Academic Publishers (The Kluwer international series in engineering and computer science; Natural language processing and machine translation, edited by Jaime Carbonell), 1994, xviii+478 pp; hardbound, ISBN 0-7923-9462-3, US $88.00, pounds 64.00, Dfl 180.00

Collective intelligence in computer-based collaboration *John B. Smith* (University of North Carolina) Hillsdale, NJ: Lawrence Erlbaum Associates (Computers, cognition, and work, edited by Gary M. Olson, Judith S. Olson, and Bill Curtis), 1994, xvi+248 pp; hardbound, ISBN 0-8058-1319-5, US $59.95; paperbound ISBN 0-8058-1320-9, US $24.50

The creative process: A computer model of storytelling and creativity *Scott R. Turner* (University of California, Los Angeles) Hillsdale, NJ: Lawrence Erlbaum Associates, 1994, ix+298 pp; hardbound ISBN 0-8058-1576-7, US $49.95

The semantics of prepositions: From mental processing to natural language processing *Cornelia Zelinsky-Wibbelt (editor)* (Universität des Saarlandes) Berlin: Mouton de Gruyter (Natural language processing series 3), 1993, vii+526 pp; hardbound, ISBN 3-11-013634-1, DM 258.00

Canadian Artificial Intelligence

Special Theme Issue

Spring/Summer 1995

"AI and the Information Superhighway"

Canadian Artificial Intelligence

Intelligence Artificielle au Canada

CSCSI/SCEIO Membership

☐ I wish to join CSCSI/SCEIO and receive Canadian Artificial Intelligence ($40.00 *Cdn./yr.)

☐ I am a student ($15.00* Cdn./yr.)
and/or

☐ I am a member of CIPS ($30.00* Cdn./yr.)

Name ______________________________

Mailing
Address ______________________________

____

____

____

Please mail your membership to:

CIPS
430 King Street West, Suite 106
Toronto, Ontario
M5V 1L5

Phone: (416) 593 - 4040
Fax: (416) 593 - 5184

For more information contact CIPS or a member of the executive.
*Includes Applicable G.S.T.

AI Advertisers Index

Applied AI .... OBC

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

Established in 1983, Applied AI Systems, Inc. (AAI) is the oldest Artificial Intelligence company in Canada. The emphasis of our business is on real world applications of cutting edge intelligent systems technology. AAI's long-term commitment to the "paying respect to science" approach is now reaping benefits. The company is international in approach, with its members constantly traveling the globe to meet other researchers and practitioners, and participating in all major conferences, workshops and symposia in related fields.

Research & Development Areas: ---Artificial Intelligent Robots --- Autonomous Systems --- Interface Agents & other "Intelligent Agents or Soft Robots" --- Evolutionary Systems & Artificial Life --- Neural Networks --- Fuzzy Logic --- Real Time Knowledge-Based Systems (RTKBS) --- Integration of fuzzy logic software with RTKBS --- Speech Recognition --- Service Provided to Clients: --- Custom System Development & Implementation --- Technical Research Reports --- Sales of SA Robots & AI Products --- Training --- Consulting ---

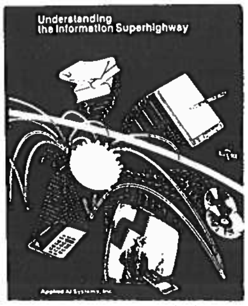

Understanding the Information Superhighway

NEW REPORT from Applied AI Systems, Inc., nearly 400-page book on the Information Superhighway: an up-to date analysis of the complexities of information highways, in addition to a detailed description of the published mergers and deals which are shaping the technology. Knowing how to navigate this highway will be crucial in government, business, education, entertainment, shopping, and banking. The report also provides a comprehensive look at leading-edge information technology, essential for discovering how the information superhighway will shape new products and business opportunities. It contains over 50 illustrations & 11 tables. $995.00 CDN

Intelligent Micro-Robots — IS Robotics — Exploring The Future

for Industrial Robot Prototyping and ALife Experiments. ISR's robots are based on the principle of Subsumption Architecture developed by Prof. Rodney Brooks of the MIT AI Lab. Now available, R3: high performance robot for ALife experiments, T1: highly maneuverable tracked robot platform, T3: T1 based human/animal interaction system, Pebbles: sealed tracked robot for hazardous environments, Genghis-II: dexterous six-legged walking robot, Nexes: advanced six-legged robot, and Hermes: durable six-legged small robot.

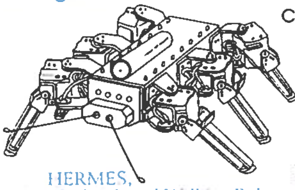

HERMES, recently developed Walking Robot

RESEARCH APPLICATIONS: --- Autonomous Agents ---- Co-operative Systems -- Exploration & Surveillance --- Dynamic Environmental Planning --- Find & Fetch Tasks --- Hardware Transport Platforms --- Evolution Experiments Using Genetic Algorithms --- Sensor Transport in Hazardous Environments, & many more.

Inexpensive Miniature Robot !

Khepera is developed by Ecole Polytechnique Federale de Lausanne in Switzerland, mostly for research institutions. Because *Khepera* is small and inexpensive, researchers can conduct experiments with multiple robots to test algorithms for distributed collective intelligence in a small area.

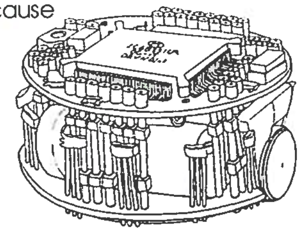

Size: Diameter 55 mm. Height 30 mm. Weight: Appx. 70 g. Processor: Motorola 68331, RAM: 256 Kbyte, ROM: 256 or 512 Kbyte. Possible applications: obstacle avoidance, wall following, target search, embodied learning algorithm, collective behavior (swarm intelligence), various Artificial Life (ALife) experiments.

--------- Small Size ---- Modularity ---- High Computational Power ---- Easy to Control ----- Programmable Highly Intelligent Operations ---- Inexpensive ---------

Khepera® Developed by EPFL lami Produced by forelec

The Premier PC Dictation Program

DragonDictate now available for Windows!

Create everything for business from complex documents to spreadsheets. Enter text, make changes, format, and print your work - entirely by voice! You can use it to control Windows functions, accessories, and applications by voice Features include: --- Direct application control and dictation --- True hands-free computing --- Excellent word recognition with minimal training --- Customizable vocabulary lets you use whatever words you want --- 120,000-word backup dictionary with speaker-independent acoustic models for quick recognition --- Dynamic adaptation adjusts to your voice and vocabulary, and environment --- Compatibility with most standard 16-bit multimedia sound cards --- Experience an easy, productive way to work in your Windows applications. Let DragonDictate put you on speaking terms with your PC!

DRAGON Dictate® *for Windows*™

For further information on pricing or ordering, please contact Applied AI Machines & Software, Inc.

Applied AI Machines & Software, Inc.
Suite 504, Gateway Business Park,
340 March Road, KANATA
Ontario, Canada K2K 2E4
Phone: + 1 613 592 7729 Fax: + 1 613 592 9762
E Mail: 73051.3521@compuserve.com

For further information on consultation or research and development, please contact Applied AI Systems, Inc.

Applied AI Systems, Inc.
Suite 500, Gateway Business Park,
340 March Road, KANATA
Ontario, Canada K2K 2E4
Phone: + 1 613 592 3030 Fax: + 1 613 592 2333
E Mail: 71021.2755@compuserve.com