Canadian Artificial Intelligence

Intelligence Artificielle au Canada

January/janvier 1989 No. 18

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

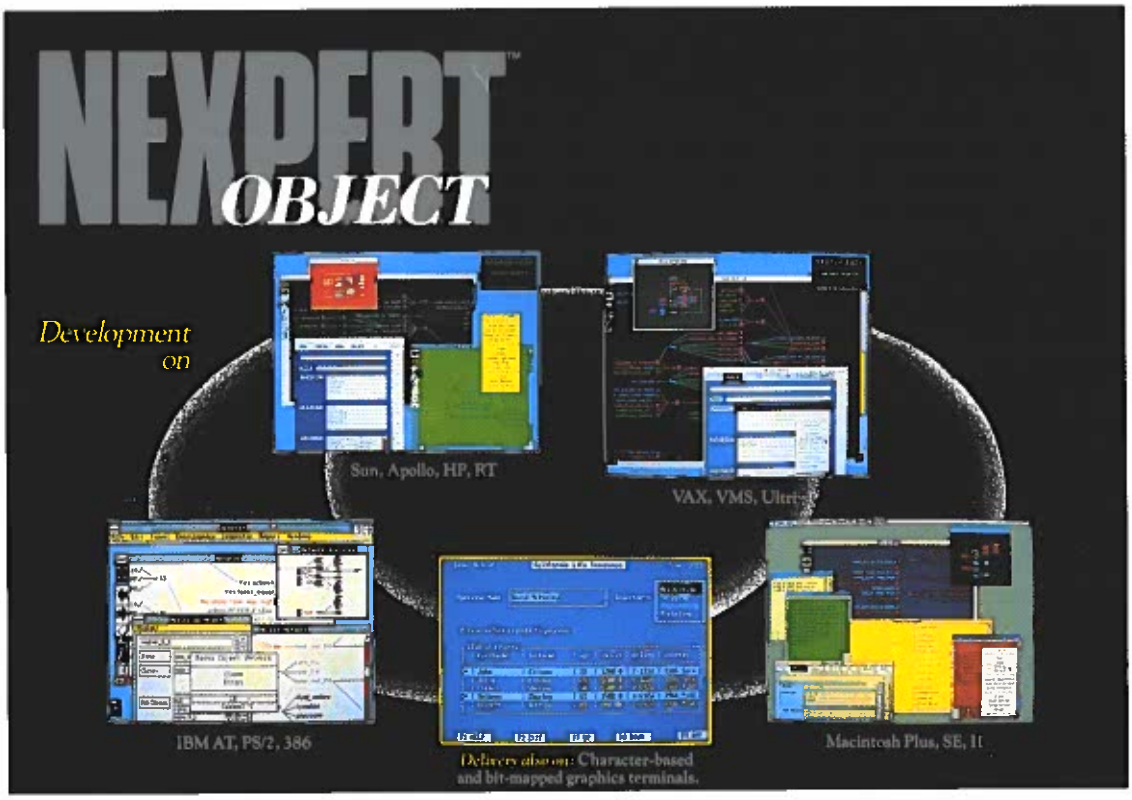

REASONS FOR NEXPERT'S SUCCESS

Links to popular conventional languages

- C
- COBOL
- FORTRAN
- Pascal
- Ada
- Assembly

Effective linkage with other databases

– provides direct system in and out calls through the Library

- Oracle
- Informix
- dBase
- Guide
- Sybase
- RdB
- Lotus 1-2-3
- Hypercard (Q1/89)
- Ingres
- SQL/DS
- Excel

NEXTRA – first truly useful Knowledge Acquisition aid

- graphically induced acquisition process with automated repertory grids
- hierarchical clustering with spatial representations
- inductive engine for example-based analysis
- graphical feedback of elicited knowledge

Linkage to other graphics packages

Ability to graphically represent macrostructure of the knowledge base

- DataViews
- Ease+
- AiVision
- PCPaint/Brush
- MACPAINT

Cross-compatibility across a wide range of platforms

– Quick and efficient adaptation to new platforms due to advanced Software Engineering techniques.

- development on VAXSTATION II, III & 2000 (VAX/VMS, ULTRIX) with delivery on all VAX computers
- Sun, HP, Apollo, and MicroVax workstations
- Hewlett-Packard 9000 series 300 and 800
- 80386 Unix workstations running X-windows
- IBM PC AT, PS/2 and compatible 80386-based PCs running Microsoft Windows
- Macintosh family (Plus, SE, II, IIx)
- IBM mainframe under VM (the 30xx, 43xx, and 9370 machines)
- Several foreign platforms (European, Japanese)

Runtime Package

- serious consideration given to runtime/delivery issues
- dual user interface concept: *end-user interface* not identical to *developer interface*
- highly automated end-user interface design tool

Phone or write for more information

Applied AI Systems, Inc.
Gateway Business Park
300 March Road, Suite 602
Kanata (Ottawa), Ontario, Canada K2K 2E2
Tel. (613) 592-3030 • Fax. (613) 592-2333

Canadian Artificial Intelligence

Intelligence Artificielle au Canada

No. 18 January/janvier 1989

Contents

Canadian Society for Computational Studies of Intelligence

Founded 1973

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

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

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

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

Fondée 1973

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

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

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

Canadian Artificial Intelligence

Intelligence Artificielle au Cana-

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

Submissions:
Canadian Artificial Intelligence is published quarterly by CSCSI/SCEIO and is a benefit of membership in the society. *Canadian AI* solicits contributions in English or French on any matter related to artificial intelligence, including: articles of general interest; descriptions of current research and courses; reports of recent conferences and workshops; announcements of forthcoming activities; calls for papers; book reviews and books for review; announcements of new AI companies and products; opinions, counterpoints, polemic, controversy; abstracts of recent publications, theses, and technical reports; humour, cartoons, artwork; advertisements (rates upon request); anything else concerned with AI. Paper or electronic submissions are welcome. Electronic submissions are preferred and should be unformatted. *Canadian AI* is published in January, April, July, and October. Material for publication is due six weeks before the start of the month of publication.

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

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

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

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

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

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

ISSN 0823-9339

Executive Notes

The standard aspiration in artificial intelligence is to create a program that will astound the user (and the researcher's funding agency) by answering questions about a children's book, by learning to discriminate a dromedary from a Bactrian camel (or, more ambitiously, an elephant from a jar of marmelade), by deciding why John left the restaurant in high dudgeon when his hamburger was burnt — or any one of the myriad tasks that AI researchers proudly undertake. The program will perform the chosen feat by gratefully accepting the input along with whatever crumbs of "knowledge" the programmer deigns to code into its innards, without so much as seeking clarification of the problem — let alone inquiring about its purpose, environmental implications, or social context. Failure to demonstrate this ability is typically attributed to the system possessing insufficient world knowledge or an inadequate inference mechanism; and steps are taken to remedy the problem by adding whatever is necessary to make the system work (and gain a PhD).

The mechanism used by the program to solve the problem is pure contemplation, deep-thought navel-gazing. The system cogitates idly (or furiously, depending on the MIPS available to it) upon the input to answer questions about it later. From the program's point of view there is no sense of purpose, no rhyme or reason, no connection with the world — it operates in a timeless, unsituated, unpopulated void. No wonder AI programs seem catatonic — they suffer from existential angst.

Navel-gazing is not known as an especially powerful problem-solving strategy for people. Even the most contemplative amongst us acknowledges the severe limitations of pure thinking, and the necessity of stimulating the creative process through interaction with the environment and with other people. On a macro scale, we strive to surround ourselves with intelligent people; attend lectures and seminars; visit colleagues; travel to conferences; communicate our ideas constantly; read the net; perform undirected, speculative experiments. On a small scale, when faced with a particular problem we try to clarify it, seek input from others, have them criticize our half-baked ideas, play around with the problem — and as often as not either conclude that it is not worth solving or turn it into a radically different problem.

Taking the lead from the theory of computational complexity, we define the class **NG** of AI problems that are tractable — in other words, can be solved in polynomial time — using navel-gazing methods. Further, we define the set **G** of problems that can be solved in polynomial time using techniques that are not limited to pure introspection. These techniques include seeking clarification, advice, and even solutions to similar problems, from an NI (naturally intelligent) system. The key question then becomes whether **G = NG**. Most current research efforts in AI are predicated on a positive answer to this question. In contrast, we believe that more attention should be given to the consequences of a negative answer to the **G = NG** hypothesis. In other words:

> There exist problems, called **NG-hard**, for which efficient (ie polynomial-time, deterministic) solutions require assistance from an NI system;

or, to put it more colloquially:

> Efficient solutions to **NG**-hard problems must employ human-computer interaction as well as AI techniques — and require some humility on the part of the machine.

Further insight can be gained by defining the class **Co-NG** of problems that can be proved insoluble using navel-gazing methods. History is littered with problems, like heavier-than-air flight, long thought to be in **Co-NG** but which have actually been solved. In fact, it seems highly plausible that the class **Co-NG** is empty — that there are no practically interesting problems that can be proved impossible to solve using navel-gazing methods alone.

On the face of it, it may seem a noble aspiration to build a system that solves problems given to it using pure reason alone, spurning human assistance. But from a commonsense point of view it appears absurd to think that any intelligent system, be it artificial or natural, can induce the

notion of a cup in all its human richness from a handful of examples, or determine the base motives of "John" as he asks "Mary" out to dinner, without some human assistance to fit events into an environmental and social context.

Although AI systems as a breed seem to suffer from serious congenital defects in their user interface, the idea of augmenting artificial intelligence by interacting with an NI system is, of course, not new. Winograd and Flores' recent book tells us about autopoietic systems: systems that are best understood as an organization of interacting autonomous and self-perpetuating agents. Our proposal is to break the tyranny of "batch AI" and make systems more interactive, more social, more humble, but ultimately more successful. Only when we can imbue computer programs with real wisdom should we attempt to tackle **NG**-complete problems. Wisdom-based systems, unfortunately, may take some time to appear. In the interim, we contend that much can be gained by moving research away from the **NG** paradigm and towards closer interaction with teachers and other people.

Roy Masrani
and Ian Witten

From the Treasurer

We have been informed by the general manager of the NRC Research Journals that the subscription rate for *Computational Intelligence* will be increased in 1989. As of January 1, 1989, CSCSI members will be paying $20 per year (up from $16) for a subscription to *Computational Intelligence*.

Jan Mulder
Treasurer

Humour

Expert Systems, Artificial Intelligence and Progressive Conservatives

Robert Spring

"Progressive conservative" is an oxymoron — the words contradict each other. I hope "artificial intelligence" is also an oxymoron — something smart shouldn't be fake. Nonetheless, since 1956 when some Dartmouth conference-goers started using the term, we have been all too eager to spread it around.

Expert systems are a product of artificial intelligence research. With their moronic oxymoron legacy, it's understandable that expert systems are misunderstood. All too often expert systems are described using phrases such as "intelligent program" and "program that solves problems usually solved by a human expert". These phrases are woefully inadequate and misleading. As Dr. Tsotsos of the University of Toronto so bluntly stated: "Definitions [of expert systems] are usually tailored so that the author can use the term to describe his own work, since it would lead to potential financial advantage." Dr. Tsotsos goes on to say that "in one sense every computer program ever written that permits branching and solves some task is an expert system. But this is really a misuse of the term and does not give sufficient credit to the original work." (*CIPS Review*, Sept/Oct 1985)

So what is an "expert system"? I'll explain with the following program written in MProlog. The example avoids anything moronic or woefully inadequate.

```
X likes Z if
        X likes Y and
        Y likes Z

bourassa likes mulroney
turner likes bourassa
mulroney likes mila
```

To use this expert system you have to know that in MProlog, you ask a question by prefacing a statement with a question mark. To begin, you could ask:

```
? turner likes bourassa
```

The expert system would then do its thing called pattern matching or "unification". The question "turner likes bourassa" exactly matches a fact that the expert system already knows, so the expert system would reply:

```
YES
```

This confirms your suspicions. Does Bourassa have any other friends? To find the answer to that question, you could ask:

```
? WHO likes bourassa
```

To understand how the expert system solves this problem you have to know that in the MProlog language, words starting with a capital letter are variables. The expert system can replace variables with anything it wants. The process is called "instantiation". In this example, the expert system would discover that by replacing "WHO" with "turner", it can match the question to a fact it knows, so the expert system would reply to the above question by saying:

```
WHO = turner
```

There are no other answers so poor Turner is alone again. The question:

```
? WHO likes mulroney
```

is more difficult to answer. Of course, the expert system would immediately discover that by replacing "WHO" with "bourassa", it can match a known fact. Therefore, the expert system would first reply:

```
WHO = bourassa
```

However, the expert system can also make use of the rule at the beginning of the program. By instantiating "X" with "turner", "Y" with "bourassa" and "Z" with "mulroney", the rule reads:

```
turner likes mulroney if
        turner likes bourassa and
        bourassa likes mulroney
```

The conditions "turner likes bourassa" and "bourassa likes mulroney" both match known facts, so the expert system learns that "turner likes mulroney". Therefore, in reply to the question about who likes Mulroney, the expert system would give two replies:

```
WHO = bourassa
WHO = turner
```

Obviously, Mulroney is a very likable guy. Obviously, expert systems lack common sense.

Can you prove that all three guys like Mila? Be sure to exhaustively instantiate and unificate.

New ISTC Programs

Three new programs; termed the cornerstone of the Department of Industry, Science and Technology; were announced in late October.

"Industry, Science and Technology Canada has been assigned the task of ensuring the international competitiveness of Canadian industry and encouraging a standard of business excellence that will secure a place for Canada in the front rank of industrial nations," according to Mr. Robert R. de Cotret, Minister of Regional and Industrial Expansion. To achieve this task, the department has initiated programs of strategic technologies, sector competitiveness initiatives, and business information and development services.

The three programs are budgeted at $400 million over four and one-half years, with the strategic technologies program accounting for about one-half of the total.

Of particular interest in the announcement of these programs is the explicit acknowledgement by the federal government that business activity is a spectrum of activities beginning with basic research, leading through applied research and product development to production and marketing. These programs are intended to change the historic emphasis on the production phase of business and to fund the components of the chain more evenly.

NATO Science Affairs Division Announces Advanced Educational Technology Program

The NATO Science Committee has introduced an advanced educational technology program "to stimulate and facilitate international collaboration among scientists of the member countries in areas related to advanced educational technologies." The intention is to help fund workshops, Advanced Study Institutes, collaborative research and academic visits.

Eight areas of interest have been selected:

- Task analysis and expert systems
- Simulation of processes
- Modeling and diagnosis of the student
- Tutorial strategies and learner control
- Interactive technologies
- Microworlds and problem solving
- Interface design
- Evaluation methodologies for innovative implementation.

Applications for funding will be considered on a competitive basis, and only a small number will be awarded each year. Priority will be given to proposals which deal with one or more of the areas of interest, which will result in interdisciplinary and innovative activities, and which will improve the interchange of ideas and professional abilities of the participants.

Further information, and applications forms, are available from Scientific Affairs Division (Advanced Educational Technology), NATO, B - 1110 Brussels, Belgium.

Nouveaux programmes au MIST

Trois nouveaux programmes; identifiés comme la pierre angulaire du ministère de l'Industrie, des Sciences et de la Technologie; ont été annoncés à la fin du mois d'octobre.

"Industrie, Sciences et Technologie Canada s'est vu confier la tâche d'assurer la compétitivité de l'industrie canadienne sur le plan international et d'encourager l'excellence commerciale en tant que norme, ce qui garantira au Canada une place au premier rang des nations industrielles.", explique M. Robert de Cotret, ministre de l'Expansion industrielle régionale. Pour accomplir cette tâche, le MIST a instauré trois nouveaux programmes: le programme de technologies stratégiques, le programme d'initiatives de compétitivité sectorielle et le programme de services d'information et de développement commercial.

Le budget de ces trois programmes s'élève, pour une période de quatre ans et demi, à 400 millions $ dont près de la moitié est consacrée au programme des technologies stratégiques.

Il est particulièrement intéressant de constater, dans l'annonce de ces trois programmes, que le gouvernement fédéral reconnaît explicitement que l'activité commerciale consiste d'un ensemble d'activités qui doit commencer par la recherche fondamentale et poursuivre, par la recherche appliquée et le développement de produit, jusqu'à la production et la commercialisation. Ces programmes visent en principe à diminuer l'importance que l'on a donnée historiquement à la phase de production et à financer chaque maillon de la chaîne plus également.

La Division des affaires scientifiques de l'OTAN annonce le Programme de technologie éducative avancée

Le Comité des sciences de l'OTAN vient de lancer un programme de technologie éducative avancée afin de "stimuler et faciliter la collaboration internationale des scientifiques des pays membres dans les domaines reliés aux technologies éducatives avancées". Le but de ce programme est d'aider au financement d'ateliers, d'instituts d'études avancées, de recherches menées en collaboration et de visites académiques.

Huit domaines d'intérêt ont été choisis:

- L'analyse des tâches et les systèmes experts
- La simulation de processus
- La modélisation et le diagnostic de l'étudiant
- Les stratégies tutorielles et le contrôle de l'apprentissage
- Les technologies interactives
- Les micromondes et la résolution de problème
- La conception d'interface
- Les méthodes d'évaluation pour une mise en application innovatrice

Les demandes de financement seront évaluées par concours, et quelques-unes seulement seront retenues chaque année, la priorité étant accordée aux propositions touchant à un ou à plusieurs des domaines d'intérêt identifiés, ce qui résultera en des activités interdisciplinaires et innovatrices et favorisera l'échange d'idées et de

R&D Funding Continues to Decline

Canadian R&D funding continues to decline on a yearly basis, and is expected to fall to 1.2% of GDP for 1988. This compares to 1.3% for 1986, 1.4% in 1984 and 1.69% in 1983. Canada now ranks last of eight major industrial countries (Britain, Canada, France, West Germany, Japan, Netherlands, Sweden, and the US) in a number of R&D metrics including R&D spending as a percentage of GDP, industrial R&D spending as a percentage of GDP, the number of scientists and engineers per capita, and the number of international patents granted per capita. The poor showing in the area of industrial R&D spending has been attributed to a branch plant economy; explanations for the other "lasts" may require more probing thought.

Wiegand Award for Canadian Excellence

Nominations are invited for the Wiegand Award for Canadian Excellence, a new award established at the University of Waterloo by the Wiegand Foundation. This award seeks to recognize Canadians (or permanent residents) who have made an outstanding contribution to our understanding of the human dimensions of science and technology or to "making technology more human." Their work will typically have the character of "widening our horizons and stirring our imagination," and candidates will have exhibited excellence to a degree meriting special recognition. The Award is administered by the University of Waterloo's Center for Society, Technology and Values, and it carries a cash value of Cdn$2500. The deadline for nominations has already passed. For further information contact: Wiegand Award Committee; University of Waterloo; Centre for Society, Technology and Values; MC 4049; Waterloo; Ontario N2L 3G1.

New Journals

The Computing Research Laboratory of New Mexico State University at Las Cruces announced two new AI journals, the *Annals of Mathematics and Artificial Intelligence*, and the *Journal of Experimental and Theoretical Artificial Intelligence (JETAI)*. Editors: Box 30001/3CRL, New Mexico State University, Las Cruces, NM 88003-0001.

The *IEEE Transactions on Knowledge and Data Engineering* will publish its first issue in March 1989. For submissions, contact Benjamin W. Wah at the University of Illinois at Urbana-Champaign. Phone: (217) 333-3516 or (217) 244-7175.
Email: wah%aquinas@ucx.cso.uiuc.edu.

AI Bibliography and Genealogy

The Computer Science and Engineering Department of the University of California at San Diego is developing a collection of bibliographic references to the literature of AI and cognitive science. A system is being developed to make this collection available over Internet. Work is now being carried out to augment the database with "a family tree of the intellectual lineage of the authors". Interested researchers should contact: Richard K. Belew at UCSD. Phone: (619) 534-2601 or (619) 534-5948.
Email: rik%cs@ucsd.edu.

compétences professionnelles chez les participants.

On peut obtenir des renseignements supplémentaires, et des formulaires de demande, à la Division des affaires scientifiques (technologie éducative avancée), OTAN, B - 1110 Bruxelles, Belgique.

Le financement de la R-D continue à décroître

Le financement de la R-D canadienne continue à décroître d'année en année et l'on prévoit qu'il tombera à 1,2% du PIB pour l'année 1988, par rapport à 1,3% pour 1986, 1,4% pour 1984 et 1,69% pour 1983. Le Canada se classe maintenant au dernier rang des 8 pays les plus industrialisés (Allemagne de l'Ouest, Canada, Etats-Unis, France, Grande-Bretagne, Japon, Pays-Bas, Suède) dans un bon nombre de statistiques reliées à la R-D dont, notamment, le pourcentage des dépenses de R-D et de R-D industrielle, le nombre de scientifiques et d'ingénieurs per capita, et le nombre de brevets internationaux accordés per capita. La pauvre performance affichée par le Canada au niveau des dépenses de R-D industrielle a été attribuée à une économie de succursale; quant aux autres "dernières places", leur explication pourrait demander une analyse plus approfondie.

Le Prix Wiegand de l'Excellence canadienne

Les candidatures sont attendues pour le Prix Wiegand de l'Excellence canadienne, un nouveau prix institué à l'Université de Waterloo par la Fondation Wiegand. Cette récompense se veut une reconnaissance des canadiens (ou des résidents permanents) qui se seront illustrés par leur éminente contribution à notre compréhension des dimensions humaines de la science et de la technologie, "à rendre la technologie plus humaine". Les candidats, par des travaux qui, typiquement, auront su élargir nos horizons et stimuler notre imagination, devront avoir démontré un niveau d'excellence méritant une reconnaissance spéciale. Le Prix est administré par le Centre for Society, Technology and Values de l'Université de Waterloo et comporte une bourse de 2500$Cdn. La date limite pour les candidatures est déjà dépassée. Pour de plus amples informations, veuillez vous adresser à: Wiegand Award Committee, University of Waterloo, Centre for Society, Technology and Values, MC 4049, Waterloo, Ontario, N2L 3G1.

Nouvelles publications

Le Laboratoire de recherche en informatique de l'Université du Nouveau-Mexique à Las Cruces a annoncé deux nouvelles publications en intelligence artificielle: *Annals of Mathematics et Journal of Experimental and Theoretical Artificial Intelligence (JETAI)*. Les rédacteurs en chef: Box 30001/3CRL, New Mexico State University, Las Cruces, NM 88003-0001.

IEEE Transactions on Knowledge and Data Engineering publiera son premier numéro en mars 1989. Pour soumissions, communiquer avec Benjamin W. Wah à l'Université de l'Illinois à Urbana-Champaign. Phone: (217) 333-3516 ou (217) 244-7175.
Email: wah%aquinas@ucx.cso.uiuc.edu.

MIRRORS/II Connectionist Simulator

A general purpose connectionist simulator has been released by the University of Maryland. This system supports an extensible high level non-procedural language, an indexed library of networks, spreading activation methods, learning methods, event parsers and handlers and a generalized event handling mechanism. The software and reference manual are free.
Email: mirrors@mimsy.umd.edu,
...!uunet!mimsy!mirrors.

Patenting Expert Systems

U.S. companies are applying for and receiving patents on expert systems software. *AIWeek* (University of Miami) interviewed patent expert Robert Greene Sterne who indicated that about 110 U.S. patents include the terms "expert systems" or "artificial intelligence". Sterne is involved with a committee to automate the U.S. patent office. As a result of the move to patents, these software developers will have to insure that their systems do not infringe on another developer's patent.

DARPA and U.S. Navy Pursue Neural Nets

Both DARPA and the U.S. Office of Naval Research (ONR) are interested in research in neural networks. DARPA has prepared a report indicating that neural nets have the potential as a fundamental component in a major weapons program. This report calls for a US$400 million program on neural nets over the next eight years. The ONR is planning a five year multi-million dollar neural network research program.

Australian Artificial Intelligence Institute

SRI International; the government of Victoria, Australia; and the Australian software company, Computer Power Group have formed an independent AI research and development group. The Australian Artificial Intelligence Institute (AAII) will be based in Victoria, Australia.

Expert Systems Business is Building Up

A new report from Ovum (London, England) says that the expert systems business is building up to a new wave of growth after "a terrible period" in 1987-88. This is a conclusion from the new report by Ovum, "Expert Systems Markets and Suppliers". The chairman of Ovum, Tim Johnson, says "There are three main reasons why we expect faster growth: the number of real applications is increasing rapidly; the technology is more mature and closer to user needs; and suppliers are being more realistic in focusing their efforts on the products with long-term potential." The report costs US$735. Contact: Ovum Ltd, 7 Rathbone St, London W1P 1AF, England.

AI in Japan

A report is available entitled "Applied Artificial Intelligence in Japan: Current Status, Key Research

Bibliographie et généalogie de l'intelligence artificielle

Le Département d'informatique et de génie de l'Université de la Californie à San Diego est à mettre au point une compilation de références bibliographiques des écrits en intelligence artificielle et en science cognitive. Ils sont en train de développer un système destiné à rendre cette compilation accessible sur Internet. Ils travaillent présentement à augmenter la banque de données d'un "arbre généalogique de la lignée intellectuelle des auteurs." Les chercheurs intéressés devrait communiquer avec: Richard K. Belew à l'UCSD.
Email: rik%cs@ucsd.edu. Phone: (619) 534-2601 ou (619) 534-5948.

Le simulateur connectionniste MIRRORS/II

Un simulateur connectionniste général a été lancé par l'Université du Maryland. Ce système comprend un langage fonctionel évolué extensible, une bibliothèque de réseaux indexée, des méthodes d'activation par propagation, des méthodes d'apprentissage, des analyseurs et des manipulateurs d'éveénements et un mécanisme généralisé de manipulation d'événements. Le logiciel et le manuel de référence sont gratuits.
Email:mirrors@mimsy.umd.edu,
...!uunet!mimsy!mirrors.

Breveter des systèmes experts

Des compagnies américaines sollicitent et reçoivent des brevets pour des logiciels de systèmes experts. *AIWeek* (Université de Miami) a interviewé l'expert en brevets Robert Greene Sterne qui a indiqué qu'environ cent dix brevets américains contiennent les termes "systèmes experts" ou "intelligence artificielle". Sterne fait partie d'un comité d'automatisation du bureau des brevets américains. Suite au mouvement en faveur des brevets, ces informaticiens devront s'assurer que leur système n'empiète pas sur le brevet d'un autre développeur.

DARPA et la Marine américaine continuent dans la voie des réseaux neuronaux

A la fois DARPA et le Bureau américain de recherche navale (ONR) s'intéressent à la recherche en réseaux neuronaux. DARPA a préparé un rapport montrant que les réseaux neuronaux ont le potentiel de servir comme composante fondamentale dans le cadre d'un programme majeur d'armement. Ce rapport propose un programme de 400 millions $US sur les réseaux neuronaux réparti sur les prochaines huit années. L'ONR projette un programme de recherche de cinq ans et de plusieurs millions de dollars sur les réseaux neuronaux.

L'Institut australien d'IA

SRI International, le gouvernement de Victoria en Australie et la compagnie de logiciel australienne Computer Power Group ont formé un groupe de recherche et développement indépendant en intelligence artificielle. L'Institut australien d'intelligence artificielle (AAII) aura son siège social à Victoria en Australie.

Performers, and Strategic Focus". The report is written by Bruce Rubinger of the Global Competitiveness Council in Boston, MA. The report costs US$300. Contact: Hemisphere Publishing Corp., 79 Madison Ave., New York, NY 10016-7892.

Neural Net Chip Developed at Bellcore

Bellcore (Livingston, NJ) researchers have developed neural network hardware that can be trained to perform an exclusive-or operation.

Products and Corporations

- The European Commission, Esprit, has approved 158 proposals of 650 submitted under the first call for Esprit 2. Among the proposals accepted are two neural computing projects.
- The British AI company Expertech Ltd. has a contract to supply its Xi Plus expert system shell to the USSR as a part of a joint venture.
- Expertech reports that its PC expert system shell Xi Plus is a component of a condition monitoring system at the Dounreay fast breeder nuclear reactor.
- Technology Applications, Inc. (Jacksonville, FL) has received a contract to produce an expert system for classification of emergencies at the Indian Point 2 nuclear power plant.
- UFA (Newton, MA) has received a contract from the U.S. Federal Aviation Administration to extend an air traffic controller training system. The system will interface an expert system with a real-time simulator.
- A Taiwanese conglomerate, the Multitech Group, is offering to pay US$1.3 million for a microcomputer software program that can successfully challenge master Go players.
- Gensym and HP have agreed to jointly develop and market a version of Gensym's real-time expert system on the HP 9000 workstation.
- Gensym (Cambridge, MA) has announced the G2 network architecture for cooperating expert systems in distributed applications. Gensym has sold 70 copies of G2, their real-time expert system, and has announced a joint agreement with its first UK purchaser (Sira) to introduce G2 into the UK.
- Digitalk (Los Angeles, CA) announced Smalltalk/V Mac, a version of Smalltalk that runs on the Apple Macintosh.
- Parc Place Systems (Palo Alto, CA) has released version 2.3 of its Macintosh Smalltalk.
- AI Ware (Cleveland, OH) announced N-NET 210 a neural network based on the Functional Link Net architecture.
- Hect-Nielsen Neurocomputers (San Diego, CA) has introduced:
  - ANZA Plus/VME, a neurocomputing coprocessor board for Sun workstations, and
  - a new neurocomputing interface for the ANZA neurocomputer that will assist non-programmers to configure and run neural networks.
- Anza Research has announced NeuralBase, a 2 700 entry neural network computer-based bibliography which runs on IBM PCs.. Cost is under US$300.
- Nestor (Providence, RI) and Datacube (Peabody, MA) announced an integration of the neural network-based Nestor Development System and Datacube's Max Vision image processing workstation.

Le commerce des systèmes experts est en croissance

Un rapport récent d'Ovum (Londres, Angleterre) affirme que le commerce des systèmes experts se dirige vers une nouvelle vague de croissance après "une mauvaise passe" en 1987-88. Telle est la conclusion d'un récent rapport par Ovum: "Expert Systems: Markets and Suppliers". Le président d'Ovum, Tim Johnson, dit: "Il existe trois raisons principales pour lesquelles nous nous attendons à une croissance plus rapide: le nombre d'applications concrètes s'accroît rapidement, les techniques sont plus mûres et plus près des besoins de l'usager et les fournisseurs sont plus réalistes en concentrant leurs efforts sur les produits qui présentent des possibilités à long terme." Le rapport coûte 735$US. Communiquer avec: Ovum Ltd, 7 Rathbone St, London WiP 1AF, England.

L'intelligence artificielle au Japon

On peut se procurer un rapport intitulé: "Applied Artificial Intelligence in Japan: Current Status, Key Research Performers, and Strategic Focus." Le rapport est écrit par Bruce Rubinger du Conseil de la compétitivité globale à Boston, MA. Le rapport coûte 300$US. Communiquer avec: Hemisphere Publishing Corp., 79 Madison Ave., New-York, NY 10016-7892.

Puce de réseaux neuronaux développée à Bellcore

Les chercheurs de Bellcore (Livingstone, NJ) ont développé un matériel de réseaux neuronaux qui peut être formé à exécuter une opération ou exclusif.

...(continued from previous column)

- The University of Southern California at Los Angeles has released a video package of 24 taped lectures on neural computing.
- Inference Corp. (Los Angeles, CA) has announced two expert system tools in the ART-IM (Automated Reasoning Tool for Information Management) product line. The products are ART-IM/MS-DOS and ART-IM/MVS.
- Ferranti (Cwmbran, Gwent) has announced price cuts of nearly one half for the Inference ART toolkit.
- IBM announced the IBM Expert System Environment, the Expert System Consultation Environment/PC, a new version of Knowledge Tool and IBM KEE (Knowledge Engineering Environment). IBM KEE was developed with IntelliCorp and runs under MVS. Additionally, IBM announces the Enhanced Common Lisp Production System for MVS.
- Aion (Palo Alto, CA) has aligned with database system manufacturer Teradata to develop an interface between Aion's mainframe expert system tool and Teradata's Database Computer.
- Neuron Data has released a new version of the Nexpert expert system shell that interfaces with databases from Oracle, Relational Technology and Sybase.
- IntelligenceWare (Los Angeles, CA) has released its Intelligence/Compiler, a product to integrate expert systems and databases.
- Kurzweil Computer Products (Cambridge, MA)

announced a new device for the visually-impaired to scan, recognize and speak printed text. The device, called the Personal Reader, is more compact and lower in price than a previous model (8.6 kg. and US$7950).

• DEC announce VAX Decision Expert, an expert system shell built on DECwindows, DEC's implementation of the X windowing system.

• DEC has signed agreements with both IntelliCorp and the Carnegie Group.

• SD Scicon is offering a special package for academics, consisting of SD-Tuner, SD-Prolog and SD-Adviser. Cost is UK£1500.

• Carnegie Group has announced Simpak and Graphpak, two new simulation and graphics components available with Knowledge Craft, its hybrid expert system toolkit.

• The Macintosh II based microExplorer, promoted by Apple and TI as an inexpensive AI applications delivery tool, is getting some good reviews. Some benchmarks have found the microExplorer to be similar in performance to TI's Explorer II dedicated Lisp machine. Costs for development systems have been reported at US$26 795 with a base system selling for US$14 995.

• Graphael is releasing version 3.2 of G-Base, and object oriented DBMS, for the micro-Explorer. Cost is US$14 000.

• TI has announced release 5.0 software which is reported to boost microExplorer performance by up to 2.5 times with certain applications.

• Expertelligence (Goleta, CA) will supply TI with the interface to the new microExplorer Toolbox, and have acquired marketing rights to PROCYON CommonLisp.

• Knowledge Garden Inc. (Nassau, NY) has announced a Video Disk Toolkit, providing control of a video disk player from a hypertext knowledge base.

• Millennium Software (Laguna Beach, CA) has released version 2.0 of HyperX, a HyperCard expert system for the Macintosh.

• Cognition Technology (Cambridge, MA) has introduced MacSMARTS Professional. This is a hypermedia based expert system for the Macintosh.

• The Montréal firm of Le Xuan, Trans, Palma, Pouraya & Associates has developed a technique called Topdown Conceptual Analysis (TCA) to assist in the development of knowledge-based systems.

• Symbolics Inc. and Apple Computer have introduced MacIvory, an Apple Macintosh II incorporating the Symbolics' Ivory LISP processor and the Genera environment. This product is sold as a tool for delivering AI applications.

• Sun Microsystems announce release 3.0 of Sun Common Lisp. The new release is indicated to have greatly increased compilation speeds.

• Sun Microsystems and Envos Corp. have announced Envos' initial release of products on Sun workstations. The products are: Medley (the Envos software development environment), Envos LOOPS, ROOMS, Flexis and Factories.

• Sun Microsystems and Schlumberger Technologies have agreed to let Sun access Schlumberger's LISP window toolkit.

• KDS Corp. has announced version 3.3 of KDS, an IBM PC frame-based expert system shell with blackboard.

• Franz Inc. (Berkeley, CA) is porting Allegro CommonLisp to MIPS-based computers. Franz has also announced that it will implement a high performance Common Lisp on Cray Research supercomputers. Of note as well is that Allegro CommonLisp is being shipped as standard software on all NeXT computers.

• Delphi (an Italian software company) has announced an implementation of Common Lisp with an implementation of CLOS (Common Lisp Object Standard).

• Bull S.A. (Paris, France) and Lucid (Menlo Park, CA) have announced a technology transfer agreement for Lucid's Common Lisp products.

• Silicon Graphics (Mountain View, CA) announced a personal 3D graphics workstation and the multi-processing IRIS POWER series.

• Logicware (New York, NY) has released its AI software running on Silicon Graphics IRIS 4D engineering workstations.

• The accounting and consulting firm Touche Ross has acquired the Knowledge-Based Systems Centre.

• Gold Hill (Cambridge, MA) announced Goldworks II. This version of the expert system shell runs on PCs, Macintosh IIs and Suns.

• Intel (Beaverton, OR) will distribute Lucid CommonLisp for Intel's iPSC/2 family of parallel processors. As well, Intel has launched a concurrent I/O hardware and software system for the IPSC/2.

• Paperback Software announced version 2 of VP Expert for MS-DOS and a Macintosh version of VP Expert.

• Information Builders Inc. announce a Macintosh and mainframe version of LevelFive.

• Sun reports a revenue of US$1 billion for the fiscal year ended June 30, almost double from the last fiscal year.

• Texas Instruments has reorganized its AI divisions with a discontinuation of the Advanced Systems Division. TI's AI operation will now be carried out from the Computer Systems division. The change is estimated to result in 100 to 200 less employees at TI.

• Texas Instruments reports earnings up 39 percent for the second quarter ended June 30.

• Tektronix reports a US$16.7 million fiscal year loss.

• Apple Computer Inc. reports a 71 percent increase in earnings during its third quarter, ended July 1.

• IBM reports revenues up 6 percent and net down 18 percent for the second quarter ended June 30.

• Unisys reports a 34 percent increase in earnings for its second quarter ended June 30.

• Symbolics (Cambridge, MA) has reported 1988 revenues of US$81 339 000 with a net loss of US$46 036 000. For the quarter ended October 2, Symbolics has reported a profit of US$415 000.

Deadline for the April issue is 15 February

COMDALE/X
The Intelligent Expert System Development Tool

Providing you with the Competitive Edge
Because we at Comdale Technologies Inc. think
You deserve the Best

Phenomenal Capabilities ...

COMDALE/X is a rule-based expert system development tool with tremendous capabilities that are extremely easy to harness; the user-friendliness of the tool is one of its greatest attributes. All rules have a simple English-like syntax, and can accommodate unlimited conditions and conclusions, intricate mathematical, logical and string expressions, and many other features. COMDALE/X runs on the PC and PS/2 under DOS, QNX and XENIX.

The developer has the capability of implementing controlled forward and backward chaining using depth-first or breadth-first search at any point in the knowledge base. Certainty factors, confidence thresholds, fuzzy sets, mutually exclusive sets, default values, meta-rules and many other functions are supported. Graphics images can be captured and displayed, access to the operating system functions is available at any time, programs written in other languages can be called, etc. A rich and powerful set of Library functions is available which allows you to: make previously fired rules available for firing again, focus the line of thought, import data, load different knowledge bases, establish goals, optimize the search strategy, and numerous other features.

Extensive Debugging and Explanation ...

The extensive debugging and explanation facilities allow you to set breakpoints, trace execution, list rules executed, ask for justification of conclusions and explanation of questions being asked, and what the system is doing, etc. The consultation session can be saved and replayed at any time with modification of previous inputs if desired, a report can be generated, information obtained and decisions made can be exported to external programs.

Documentation ...

The documentation is complete. No prior knowledge of expert systems development is assumed, the theoretical foundation of the tool is outlined in concise terms, a wide range of examples are provided, the tool is supplied with sample knowledge bases and a tutorial.

Our Services and Support are Comprehensive and Continuous ...

At Comdale Technologies Inc. our staff of dedicated specialists can assist you in identifying and defining applications, developing and delivering customized knowledge bases according to your specifications, and last but by no means least, we provide training. Because we know that it is important to you to have someone on the inside who is familiar with this leading edge technology, we offer the following courses:

- Executive seminar on Expert Systems
- Expert Systems training course
- Knowledge Engineering training course.

The ultimate goal of COMDALE/X is to increase productivity at a lower cost, that's COMDALE's success.

Ask us about COMDALE/C, our real-time expert system development tool for process monitoring and control.

Comdale Technologies Inc.
833 The Queensway, Suite 202
Toronto, Ontario
Canada, M8Z 5Z1
Tel (416) 252-2424

CIAR and Artificial Intelligence in Canada: A Story of Insiders and Outsiders

by Connie Bryson

L'ICRA et l'IA au Canada: une histoire de gens "en dedans" et de gens "en dehors"

RÉSUMÉ: L'Institut canadien des recherches avancées (ICRA), un réseau plutôt qu'une institution de briques et de mortier, a pour mission de mettre sur pied des programmes dans des domaines de recherche qui représentent des défis intellectuels majeurs, qui peuvent contribuer à la formation de "forces" en recherche au Canada et qui requièrent l'implication d'un institut national et un mécanisme de réseau pour établir une base canadienne solide. Il existe cependant un problème fondamental qui préoccupe plusieurs membres et aussi non-membres de l'ICRA: la fragmentation de la communauté de l'IA entre gens "en dedans" et gens "en dehors". Le plus grand défi peut-être, résultant de l'établissement du programme Intelligence artificielle et robotique de l'ICRA ne s'adresse pas à l'ICRA mais plutôt à toute la communauté canadienne de l'IA. Il consiste à s'assurer que les "en dehors" de l'ICRA ne soient pas considérés, et ne deviennent donc pas, des chercheurs de seconde classe aux yeux des agences de financement, des gouvernements et de leurs collègues du monde entier.

The Artificial Intelligence and Robotics (AIR) program of the Canadian Institute for Advanced Research (CIAR) marked its fifth birthday this year. There was cause for much celebration.

An international review panel gave the program top marks and one only has to read the comments of CIAR fellows and scholars — excerpted from documentation prepared for the review panel — to get a sense of the undeniably enthusiastic support for the program.

"The main attraction of the CIAR post was the greater freedom to do research on topics of my choice in a university setting, free from an undergraduate teaching load. This opportunity has already turned out to be a very valuable experience. I have initiated an exciting research project of a larger scale than I have been able to undertake before, and I have had some valuable interactions with CIAR fellows...

"Already the program has attracted the attention of AI researchers that I know in the U.S. and in Europe, and I think that this kind of reputation will help us attract good students and new faculty. The program has also stirred up enthusiasm, I think, just by making Canada known as a place where there is good AI research."

"The cross-disciplinary emphasis of the program has made me much more aware of developments in neighbouring disciplines...I have sought out program members to help me solve theoretical and practical problems. Without the opportunities for interaction, that would not have happened to nearly the same extent.

"With assured CIAR support, I have become less conservative — more willing to take strategic long-term risks...Those risks are paying off. We are now in a position to exploit these new basic and applied research opportunities with the essential support of the CIAR."

"The major change in my research over the past three years has been primarily in my ability to take stock of where I want to go and to mount projects or establish collaborations that will get me there. The fellowship has also given me the time to cultivate a higher profile internationally, through invited lectures and panels at conferences..."

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

Despite the rave reviews, there is one fundamental concern shared by many CIAR members and non-members alike — the fragmentation of the AI community into "insiders" and "outsiders". Perhaps the greatest challenge resulting from the establishment of the AIR program does not face CIAR but instead faces the Canadian AI community. It must now ensure that CIAR outsiders are not considered, and hence do not become, second-class researchers in the eyes of funding agencies, governments and colleagues throughout the world.

The "Institute Without Walls"

The idea for CIAR emerged from a committee of distinguished Canadian researchers and administrators, established in 1979, to study the potential for an institute for advanced study in Canada. Its report concluded that a national institute, not affiliated with any specific institution, was required.

Sufficient support was secured from individual corporations, foundations and the Ontario government to initiate CIAR on July 1, 1982. The institute is governed by a 19-member Board of Directors, primarily made up of business leaders. The President is advised by the 26-member Research Council whose members come from industry and academia. The federal government recently agreed to match private sector support for CIAR up to $7 million over the next four years.

CIAR is a network, rather than a bricks and mortar institution, an "advanced research institute without walls", as CIAR president Dr. Fraser Mustard often calls it.

CIAR's aim is to develop programs in areas of research that:

- involve major intellectual challenges
- can build up research strengths in Canada
- require the involvement of a national institute and a network mechanism to achieve a strong Canadian base.

Five programs are currently underway: artificial intelligence and robotics, cosmology, population health, evolutionary biology, and superconductivity. A full-scale review of program direction and membership is done in the fifth year of a program's operations.

Appointments to the research programs are made in

Chairman:	Barrie J. Frost Department of Psychology Queen's University	
Members:		
Michael Arbib Department of Computer Science University of Southern California	Pierre Bélanger Dean, Faculty of Engineering McGill University	Arthur N. Bourns Professor Emeritus McMaster University
Michael Brady Chair in AI & Robotics Oxford University	Eric Manning Dean, Faculty of Engineering University of Victoria	J. Ron McCullough Vice President, Corporate Planning Spar Aerospace Ltd.
ex offico: Gordon MacNabb President and CEO PRECARN Associates Inc.	Alan Newell Department of Computer Science Carnegie-Mellon University	C. Raymond Perrault Director, Artificial Intelligence Center SRI International
Zenon W. Pylyshyn Centre for Cognitive Science University of Western Ontario	Whitman Richards Dept. of Brain & Cognitive Sciences Massachusetts Institute of Technology	

Table 1. AIR Advisory Committee

three categories:

- Fellows. Outstanding scientists or scholars in, or entering, the most productive stage of their careers. They receive full salary support for terms of five years.
- Associates. Chosen according to similar criteria as fellows, associates participate fully in the program's seminars, workshops and other meetings but do not receive salary support. The terms of associates are variable.
- Scholars. These are highly promising young researchers, generally one or two years beyond their Ph.D., whose intellectual development could benefit from a close association with the program's fellows. Scholars receive full salary support for terms of three to five years.

Candidates for appointment are identified by CIAR task forces, fellows, university officials and other knowledgeable people. Assessments of a candidate's record, abilities and potential from internationally recognized experts are studied by the appropriate program Advisory Committee.

Associates and scholars are recommended directly to the President. Fellows must be approved by the Research Council before an appointment can be made.

The AIR Program: Who's Who and What's What

In 1984, CIAR established its first program, the Artificial Intelligence and Robotics (AIR) program. Its goals are:

- To promote individual research excellence.
- To promote the development of Centres of Excellence.
- To promote the development of a national community of interacting researchers crossing regional, disciplinary, and basic-applied boundaries.
- To promote industrial research and development and the transfer of research findings from academic laboratories to Canadian industry.

Within the AIR program, areas of expertise are human and computer vision, knowledge representation and reasoning, and robotics. Program development is guided by an Advisory Committee whose members are listed in Table 1.

The 37 scientists (15 fellows, 18 associates, four scholars) currently in the AIR program come from a number of disciplines, including psychology, computer science, engineering, philosophy and neuroscience. Four universities are designated as nodes for the program — McGill, Toronto, British Columbia and Western Ontario. Program leader is Zenon Pylyshyn, a cognitive psychologist at the University of Western Ontario.

In addition to a regular yearly meeting of all AIR program members, there have been a few special workshops. These include a 1986 workshop put on by the Dalhousie neuroscience group and a "Computational Approaches to Problems in Human Vision" workshop also held in 1986. The proceedings of this workshop have been published in a book, "Computational Problems in Human Vision", volume one of a series called "The CIAR Series in Artificial Intelligence and Robotics". The second volume contains proceedings of the 1987 CIAR-sponsored international workshop on "Vision and Action: The Control of Grasping".

CIAR has also been encouraging cross-disciplinary exchanges among students. Two AI workshops have been organized by and for graduate students, particularly those of CIAR fellows and associates.

The AIR Program: The "Innies" and the "Outies"

In May, 1988, the AIR program underwent its five-year review by a panel of experts (see Table 2). The committee reviewed documentation prepared by program participants — including curriculum vitae, self-assessments by fellows, associates and scholars and a selected paper from each fellow and scholar — and made a series of site visits to UBC, McGill, and U of Toronto.

The first three paragraphs from the "Findings" section of the review committee's report sum up their findings:

"The committee concluded that the CIAR had been extraordinarily successful in promoting the development of three "world-class" laboratories in artificial intelligence at McGill University, the University of Toronto and the

University of British Columbia. The breadth of disciplines within the field of artificial intelligence makes any attempt at serial ranking impractical, but, in using the term 'world-class' the committee means 'within the top 20-30 centres of activity in the world'. In practical terms, this means that specific foci, such as the computer vision activity, and knowledge representation are amongst the best ten or so internationally...

"The fact that there are three world-class centres in a country the size of Canada, and the evidence that these centres have attained their stature in such a short period of time and with rather limited funding made a strong impression on the committee. By any objective measure the AIR program must be rated as an unequivocal success.

"The committee is convinced that the three nodes would not enjoy such a status without the initiative and support of the CIAR. The evidence was overwhelming that the CIAR fellowship program had permitted the retention or attraction of many of the key individuals needed to make the nodes viable according to international criteria..."

From these comments it is obvious that the review committee found that the AIR program had accomplished at least two of its goals — promotion of individual research excellence and development of centres of excellence.

CIAR fellowships have been key to recruiting to Canada top talent such as Hector Levesque, Edward Stabler, Geoffrey Hinton, David Lowe and Wolfgang Bibel. Perhaps even more important is the role CIAR has played in keeping talent in Canada.

"Everyone in CIAR has had very tempting offers elsewhere," says AIR program leader Zenon Pylyshyn. "But the teaching relief and assistance in networking offered by CIAR are 'perks' that keep people here."

"Before CIAR, each of the three major nodes was fragile," continues CIAR fellow and UBC professor Alan Mackworth. "Many people don't recognize this. There was great pressure from the U.S. The centres would not have achieved what they have without the CIAR. They may even have folded."

But it is precisely the AIR program's success in attracting researchers and building centres that have put it under close scrutiny. The most hotly debated issue appears to be the CIAR selection process for fellows and the resulting exclusion of one or two prominent Canadian researchers from the AIR program.

As described in the "AIR Program: Who's Who and What's What" section, members are chosen on the basis of their international visibility and research ability.

"We select people who fit into the program," says Pylyshyn. "That means looking at: where they're located, an important criterion in the early years that has been downgraded somewhat now; their area of expertise, there are three areas — vision, knowledge representation, and robotics; and the extent to which the person might naturally interact with others in the program.

"We have people in the program who are not in core AI — neuroscience and engineering for example — but they show potential for building up fruitful interactions with other members."

Mackworth says the CIAR selection system is open to the effects of "alliances, politics and friendships" but believes there is "an attempt on everyone's part to be as fair as possible".

Pylyshyn says the system "leads to a certain amount of arbitrariness" but added that CIAR is trying to "minimize this arbitrariness. We're re-evaluating the whole process. We've had the review committee input and a healthy amount of self-criticism."

Both Pylyshyn and Mackworth argue that although the selection system is imperfect, it does not mean the entire organization is flawed. They believe the problems will be remedied. According to Pylyshyn, there is a top-level commitment to change.

"You've got to understand, CIAR is an experiment, a novel idea," he says. "At the start, a number of things were not done optimally. There's a mystique about CIAR that I don't believe was intended. We're learning from our mistakes."

Pylyshyn says the selection system is now being used in a more coordinated fashion to build a focused research program.

"We're not a kind of 'royal society' that rewards people for their good work," he says. "We're attempting to build a coherent research program. We're still a long way from coherence but we're moving in that direction."

"An AI institute could have been built and it might have been more effective in terms of a coherent research program," adds Mackworth. "But in the Canadian context it would have been devastating. Our universities would be sucked dry."

Mackworth believes that some of the grumbling over the management of CIAR comes from a basic misunderstanding of the organization.

"CIAR is an elitist organization and that's what it was set up to be," he explains. "Canadian academics are used to a much different system — NSERC. NSERC is very fair and very democratic and it's gained a lot of respect for those reasons. We've come to expect that every funding organization will operate in the same way.

"But there's a downside to the way NSERC operates. There's enormous inertia. It doesn't respond well to the

Chairman:	John C. Madden Principal, STC Enterprises Vancouver, British Columbia	Secretary:	Stuart Taylor Director of Program Coordination CIAR
Members:			
Michael Arbib Department of Computer Science University of Southern California		Woody Bledsoe Department of Computer Science University of Texas, Austin	Michael Brady Chair in AI & Robotics Oxford University
Alan Newell Department of Computer Science Carnegie-Mellon University		Whitman Richards Dept. of Brain & Cognitive Sciences Mass. Institute of Technology	Boris Stoicheff Department of Physics University of Toronto

Table 2. AIR Review Committee

needs of a new science like computer science. CIAR can act quickly because it's a totally different kind of animal — a private institute that makes its own rules and decisions. It's opportunistic and undemocratic but it can do things that NSERC can't."

Another criticism levelled at the AIR program is related to its success in building up centres of excellence. Some people feel that the AIR program nodes have concentrated effort to the detriment of AI programs at universities that are not nodes.

Mackworth agrees that there is a concentration of effort but notes, "In order to get first class research you have to have critical mass. Unless you're a pure theoretician, you can't be on your own."

Randy Goebel, University of Alberta computer science professor and CIAR "outsider", says, "Initially we thought there would be a drain of people to the nodes, but it really hasn't turned out that way. On the whole I'd say the benefits of CIAR have been greater than the detriments."

By all accounts, the AIR program has had the least success with its goals of networking and developing industrial contacts. Industrial liaison seems to have fallen entirely on the shoulders of PRECARN Associates Ltd., the 35-company consortium formed to support precompetitive research in robotics and AI.

But there is still a CIAR connection. PRECARN's president Gordon McNabb credits CIAR for making the consortium possible. "CIAR established the knowledge generating base within the universities. PRECARN is based on having this competence in Canada," he says.

Networking among members has been slow to establish. "Research collaboration takes a long time to develop and lab-oriented research is unrealistic to do across thousands of miles," says Mackworth. "We're still at the stage of finding out what's realistic and what's not in terms of collaboration."

Mackworth says there have been benefits to the networks other than physical collaboration. "With CIAR, I'm totally exposed to the research of other fellows. I don't have time to go to neurophysiology conferences to find out what's important in that field. But at our meetings, everyone presents their work. It's an efficient way of getting up to speed," he says.

But CIAR meetings are a controversial subject among non-members because outsiders, unless invited, are not allowed in.

"People I've written papers with go to CIAR seminars and yet I'm not invited," says Goebel. "It's utterly frustrating."

On the other hand, CIAR members praise the small, congenial atmosphere of the meetings. There appears to be no middle ground.

Not included as a goal of the AIR program, but a benefit nonetheless, is the leveraging effect CIAR fellowships have had, particularly at the nodes. Most universities have used the salary money freed up by fellowships to hire top-level teachers, research scientists, or support staff. Benefits go to CIAR members and non-members alike.

Recognition as a CIAR node helped the AI group at the University of Toronto win the largest NSERC grant ever awarded for AI research. Being a CIAR node also helped McGill to acquire provincial government funding to establish the McGill Research Centre for Intelligent Machines (McRCIM). McRCIM now receives about $2 million in research funds and contracts and has 16 staff members and over 90 graduate students.

Although clearly a bonus for people at the nodes, there is no such leveraging effect for AI research groups at other Canadian universities. If CIAR membership means a greater likelihood of winning other grants, is it a case of the rich getting richer and the poor getting poorer?

"CIAR has done a great job of raising the awareness of AI in general in Canada, no one can fault the organization on that," says Goebel. "But when they start talking specifics, when the CIAR's 'anointed group' are referred to as the best in Canada, it makes the rest of us look like second class citizens. As a result the people outside CIAR are done a disservice."

But is this really CIAR's fault? It is a private company with its own stated goals; improving the lot of all AI researchers in Canada is not in the AIR program's mandate. Invoking the fact that CIAR receives government funding is not reason enough for it to extend its mandate. Many other private companies receive government grants and do not have to alter their corporate mission.

But at the same time, CIAR does not operate in a vacuum. The 37 AIR program members couldn't function without the rest of the Canadian AI community and as a result, CIAR's actions have implications for non-members as well as members. Goebel's worry about the meaning of exclusion from the "club" is a fundamental concern. It underlies all other criticisms about the selection process, the success of networking, and the closed conferences.

"I'm very sensitive to this criticism," says Mackworth. "The AI community in Canada is very important to me. I don't want two different camps — the 'innies' and the 'outies'."

Both Mackworth and Pylyshyn see the federal government's new Networks of Centres of Excellence program as a way to mitigate the innie/outie problem. The Institute for Robotic and Intelligent Systems (IRIS) proposal, administered by PRECARN, includes about 90 AI researchers across Canada — AIR fellows and many more people who are not members of CIAR.

"If the IRIS proposal is successful, I think it will change the AIR program. There will be less pressure on CIAR to be all things to all people," says Mackworth.

"When CIAR appeared on the scene, many of us thought our problems were over. Well, there just wasn't enough time or money to do that. The fact is that, for the most part, AI researchers in Canada are atrociously funded. The scene is dismal. If this changes, CIAR can get back to its role of supporting people doing basic research."

But even if the IRIS proposal is not successful (238 letters of intent were filed and only about 15 networks will be funded), the initiative is still a valid one. Perhaps it is time for the outies to ask the innies, who appear to have the ear of governments, funding agencies and industry, for some help in lobbying for research support.

CIAR documents cite the importance of the involvement of non-members so the stage has already been set. Such a collaboration would go a long way to answer much of the criticism of CIAR and do away with the innie/outie distinction, a fragmentation the Canadian AI community does not need.

Acknowledgments

The author thanks Zenon Pylyshyn for providing AIR program documentation prepared for the five-year review.

Les aléas d'un laboratoire d'intelligence artificielle

par Philippe Duchastel

The Risks of an Artificial Intelligence Laboratory

ABSTRACT: This article presents my views about the problems and the satisfactions we encountered during the first few years of operation of our university-based AI laboratory.

Introduction

Le Laboratoire d'Intelligence Artificielle en Education (LIAE) fut créé à l'Université Laval, à Québec, au début de 1986 et il possède donc aujourd'hui une expérience de quelques deux ans et demi. Le LIAE a ceci de particulier, qu'il fut créé par un chercheur non-informaticien et qu'il fut créé dans une faculté des sciences de l'éducation. Comme nous le verrons plus loin, ces facteurs ont été source de certains problèmes fondamentaux.

Le but de cet article est en fait de présenter les problèmes que nous avons dû affronter au cours de ces dernières années, et qui sont encore avec nous sous une forme ou une autre. Si j'expose ces problèmes d'une facon ouverte, c'est dans un double but: d'une part, sensibiliser certains secteurs de l'appareil de recherche en intelligence artificielle à l'existence de ces problèmes, et d'autre part faire profiter de notre propre expérience (en autant que cela est possible) les autres groupes de recherche au pays qui se constituent en laboratoire. Il est à reconnaître bien sûr que tout laboratoire a sa propre expérience particulière, laquelle est peut-être bien peu généralisable; néanmoins, il est possible que d'autres retrouvent ici un certain reflet d'une expérience particulière qu'ils vivent ou qu'ils ont vécue dans leur contexte à eux.

Afin d'éviter un ton déplaisant à cet article, je vais également présenter, lorsqu'approprié, les aspects positifs de notre expérience à date. Je vais donc en fait traiter, dans ce qui suit, des succès et des échecs de notre laboratoire.

Création du Laboratoire

Comme probablement beaucoup de choses dans la vie, la création du LIAE prit sa source dans une circonstance fortuite. En effet, c'est au hasard d'une conversation pendant que nous nous rendions à une réunion dans une ville éloignée que le doyen de la Faculté des Sciences de l'Education, à l'Université Laval, prit intérêt à développer à la Faculté une expertise dans une technologie émergente, celle de l'enseignement intelligemment assisté par ordinateur (EIAO). J'allais bientôt quitter mon poste dans une institution voisine, et le doyen me proposa donc de venir à Laval pour y développer ce secteur de recherche. Ce que je fis, avec enthousiasme et beaucoup d'aspirations.

Très tôt, je vis l'utilité de constituer une structure qui à la fois nous donnerait une certaine visibilité, tant à l'interne qu'à l'externe, et nous conférerait une certaine réputation, initiallement par pure association à cette mystique qu'est l'intelligence artificielle, mais éventuellement par nos propres activités de recherche qui y seraient encadrées. Je m'associai donc deux collègues intéressés à l'EIAO et, avec le soutien de notre directeur départemental, nous installames une affiche à la porte de notre "labo": le LIAE. Le LIAE était né, et cela sans autre formalité.

Le LIAE a toujours été une structure informelle, c'est-à-dire sans reconnaissance formelle des autorités universitaires. Cela lui évite la bureaucratie qui est inévitablement rattachée à la formalité (et lui confère donc beaucoup de flexibilité), tout en lui accordant une bonne part du prestige qui résulte des liaisons externes d'un tel laboratoire. D'ailleurs, nos premières actions ont été dans le domaine des relations externes: instituer une série de documents de recherche pour disséminer nos recherches, et nous faire connaître de la communauté canadienne (une description du LIAE et de ses activités parut dans *Intelligence Artificielle au Canada* en juin 1986).

Aspirations et difficultés

Le LIAE a pour mission de promouvoir à l'université (et en particulier à la Faculté des Sciences de l'Education) la recherche en intelligence artificielle appliquée à l'éducation, domaine qui comprend l'EIAO et les systèmes experts reliés à l'éducation. A l'époque de sa création, le LIAE se donnait également pour mission globale de favoriser le développement d'une expertise québecoise et canadienne dans l'application de systèmes intelligents en éducation (en 1986, aucun laboratoire spécialisé n'existait dans ce domaine d'application, bien que certains chercheurs chevronnés tels McCalla et Jones y étaient très actifs).

L'action du LIAE était, et est encore, orientée vers le développement: sa programmation est centrée sur le développement de prototypes, plutôt que d'être orientée vers des analyses théoriques, quoique cette dernière perspective est nécessairement impliquée dans tout développement qui se veut être à la pointe de la recherche scientifique en IA. A date, plusieurs projets de prototypes de systèmes tutoriels intelligents ont été entrepris: un en géographie, un en nutrition, un en optique, et d'autres encore sont au stade de conception; une maquette de système expert en design pédagogique a également été réalisée. Ces divers prototypes ont été réalisés en Lisp, en Prolog, ou encore en HyperTalk.

Les problématiques que nous abordons par la construction de ces prototypes sont principalement le dialogue personne-machine, l'initiative de l'étudiante en interaction avec un système riche et flexible, les stratégies tutorielles, en particulier concernant la motivation, le potentiel du "hypermedia", et la correction de fausses conceptions scientifiques.

Cependant, aucun de ces prototypes n'a évolué vers un système pratique ou complet, non pas parce qu'il n'y a pas intérêt à le faire, mais bien parce que nous n'avons pu obtenir aucun financement pour le faire.

En trois années de demandes de subventions, aucunes de nos propositions de recherche n'a su intéresser suffisamment nos pairs pour nous placer dans la catégorie de chercheurs subventionnés. Il va sans dire que cette

Philippe Duchastel est le directeur du Laboratoire d'Intelligence Artificielle en Education à l'Université Laval.

situation fort embarassante nous cause et des soucis, et des ennuis très pratiques.

En regardant de nouveau cette dizaine de propositions de recherche que nous avons adressées aux fonds subventionnaires, et en examinant les remarques des évaluateurs qui nous furent éventuellement communiquées, on peut déceler une certaine inexpérience bien relative de notre part comme un des facteurs ayant joué un rôle dans le résultat: nous étions nouveaux dans le domaine et notre perspective était ambitieuse (malgré que, même aujourd'hui, elle ne paraît pas l'être trop).

Quoiqu'il en soit pour cela, un second facteur qui semble avoir joué un rôle dans le résultat est l'interdisciplinarité du domaine et notre propre unidisciplinarité nominale. Cela joue de plusieurs façons. D'abord sur le plan méthodologique: nos collègues en éducation apprécient peu la méthodologie de prototypage qui est centrale à la recherche en intelligence artificielle. Pour eux, cela ressemble sûrement trop à du développement de matériel pédagogique et ne constitue donc pas en soi de la recherche, même appliquée. Le prototypage est une méthodologie trop floue pour des gens qui recherchent des designs expérimentaux et des données statistiques pour valider des hypothèses de recherche. Une méthodologie informatique dans des comités d'éducation fut donc une source de difficultés de nature interdisciplinaire.

Ensuite, il y a l'habileté jugée des chercheurs. Au tout début, j'avais adressé une demande au CRSNG, car c'était après tout pour le développement d'un prototype informatique; mais aussitôt que mon CV fut examiné, la demande fut transférée au CRSH, à un comité d'éducation. A plusieurs reprises également, on nous reprocha de vouloir faire de l'informatique sans être informaticiens, même après que nous ayons réalisé des prototypes fonctionels en Lisp et en Prolog, et malgré une expérience en informatique appliquée qui remonte à plus de 15 ans. Voyant bien la nécessité de la chose, nous avons cherché à nous associer des collègues informaticiens, mais sans grand succès hélas (l'informatique est vaste et chacun a ses propres projets à faire avancer).

Nous nous retrouvons donc coincés entre deux disciplines, en train de développer des systèmes informatiques en éducation, alors que chacune de ces disciplines nous considère un peu bizarre pour cela. Et certes un peu trop bizarre pour nous alouer des fonds de recherche, ces fonds étant rares et en grande demande; d'autant plus que le budget du CRSH est près de cinq fois inférieur à celui du CRSNG. Et les dépenses en R&D au Canada en intelligence artificielle sont proportionnellement de loin inférieures à celles des autres pays comparables. Il y a trop peu d'argent, donc, pour soutenir des projets qui sortent quelque peu de la norme traditionnelle.

Il n'y a pas que l'embarras d'espoirs à reconsidérer constamment qui nous cause soucis. En effet, le manque de succès auprès des organismes externes est en train de miner la confiance interne en nos projets de la part des autorités facultaires et universitaires. Nous comptons sur cette confiance pour les quelques fonds de soutien internes qui font marcher le laboratoire à l'heure actuelle, mais nous la voyons continuellement s'éroder avec le temps.

Ces difficultés, qui sont à la base dues à l'incapacité du Canada de subvenir financièrement aux aspirations de recherche de plusieurs de ses chercheurs, nous font parfois rêver à des possibilités d'avantage propices ailleurs.

Cependant, le laboratoire a jusqu'à date survécu à ces difficultés et il continu à solliciter des fonds pour mener à terme ses recherches. Incidemment, nous sommes portés de plus en plus à nous tourner vers des sources potentielles de financement qui sont moins traditionnelles (celles-ci nous ayant tourné le dos), plus informelles et davantage politiques, c.-à-d. où les connaissances personnelles des gens jouent un certain rôle. Ce n'est certes pas là une situation que nous préférons, mais elle semble davantage réaliste dans notre situation que celle que nous avons poursuivie à date.

Côté matériel informatique, nous avons obtenu la confiance de la compagnie IBM-Canada, qui nous a généreusement offert un PC RT sur lequel nous avons continué notre développement de GEO en Lisp. La Fondation Apple Canada, de son côté, nous annonçait en novembre 1987 l'octroi d'un Mac II augmenté, mais les embêtements bureaucratiques qui suivirent ont fait que nous n'avons toujours pas vu la couleur de cet appareil et nous nous demandons vraiment s'il sera de fait livré un jour (peut-être un an après l'annonce de l'octroi?).

Malgré nos ressources limitées, nous avons pu poursuivre nos recherches au fil des dernières années, et ainsi présenter plusieurs communications à des réunions savantes (grâce incidemment au soutien du CRSH et du MRIQ québecois) et publier plusieurs articles dans des livres et des revues, de même que participer à plusieurs comités d'experts dans le domaine. Même si nos systèmes informatiques demeurent incomplets, les prototypes que nous avons réalisés sont une source de satisfaction, de même qu'une base sur laquelle nous pouvons entreprendre de nouvelles aventures de recherche, d'autant plus qu'un nombre toujours croissant d'étudiants et d'étudiantes s'intéressent à ce domaine.

Conclusion

La vie d'un laboratoire de recherche en intelligence artificielle est certes excitante. Le tout demeure cependant assez fragile, surtout lorsque situé dans un contexte interdisciplinaire.

De quelle façon notre expérience particulière (et il faut bien dire qu'elle l'est) rejoint celle de d'autres équipes de recherche au Canada reste à déterminer. Notre contexte en est un qui est universitaire, et il est sûr qu'un laboratoire dans un contexte industriel jongle avec des difficultés autres. Néanmoins, nous croyons que notre expérience reflète une certaine réalité dans la structure de la recherche, et dans ce domaine et au Canada, et c'est pour partager cette expérience que ces propos sont présentés. Il va sans dire qu'il s'agit de propos personnels qui relèvent de ma propre perspective des choses.

L'échéance pour le numéro d'avril est le 15 février

Artificial Intelligence Research at CRIM

by Jennifer Muise and Diane Goupil

CRIM (Computer Research Institute of Montreal) has designated artificial intelligence as a research priority. The other aspects of computer research at CRIM reflect the central theme; often, several research domains are involved in a single project. Other fields in which CRIM is involved include software engineering, man-machine interfaces, computer communications, CAD/CAM and VLSI. CRIM's laboratories offer a favorable environment for the research and development of projects.

Expert Systems

We are exploring various aspects of knowledge representation and reasoning methods, as well as software engineering methods for expert systems design and development, and application-specific problems. Expert systems and knowledge-based systems are a major area of research and development at CRIM. We apply expert systems and knowledge-based systems to various domains in various projects.

Diagnosis of problems is studied in a PRECARN project with Hydro-Québec and Alcan. A "family" of expert systems will perform diagnoses on a set of electrical equipment, in terms of both maintenance and repair. The project involves a knowledge acquisition model for the design of a general diagnostic expert system for different types of electrical equipment. This project is under the scientific supervision of Drs. De Mori, Galiana and Lowther of McGill University.

The modeling of knowledge was the focus of a project with the CGI Group supervised by Dr. Vaucher of the Université de Montréal. An intelligent prototype of a data modeling expert system was developed.

The solving of problems caused by pitch deposits in kraft mills is being addressed in a project with PAPRICAN (the Pulp and Paper Research Institute of Canada). The expertise of Dr. Larry Allen, PAPRICAN's pitch expert, is being incorporated into an expert system which solves these problems.

How may an expert system produce and explain meteorological forecasts? This is the topic of a project for Canadian airports under way with Météoglobe Canada. A prototype capable of predicting low cloud cover has already been developed. This project is under the scientific supervision of Drs. Newborn, of McGill University, and Zwack, of the Université du Québec à Montréal.

A system to assist in verification and troubleshooting of flight simulation models is the object of a joint study with CAE Electronics. The principal goal is to design a knowledge-based system which will assist engineers with the aerodynamic models of airplanes for simulation. Drs. De Mori and Bélanger of McGill University provide the scientific direction for this project.

Jennifer Muise and Diane Goupil are research associates at the Centre de recherche informatique de Montréal (Computer Research Institute of Montreal).

Speech Recognition

Many industries are taking note of speech recognition and its applications. Dr. De Mori has introduced this domain to CRIM in several areas. Several approaches have been explored, including that of neural networks.

A vocal interface for an existing expert system is the subject of a joint project with IBM. The expert system will ask a series of multiple choice questions, and the interface will allow the user to respond vocally, rather than through the keyboard. Recognition of rapid speech of air traffic controllers is being explored in a project with CAE Electronics.

Robotics

An intelligent robot is the subject of a research project subsidized by PRECARN, along with the Institut de recherche d'Hydro-Québec (IREQ) and the McGill Research Centre for Intelligent Machines (McRCIM). The robot will be guided remotely by an operator and will automate the maintenance of high tension lines. The robot will have its own knowledge base which will permit it to perform certain tasks automatically and to consult the operator on its decisions. This project is under the supervision of Drs. Daneshmend, Hayward, Levine and Zucker of McGill University.

Natural Language Understanding

Montreal has had a rich community of researchers in natural language understanding since the 1970's with the TAUM project in automated translation which took place at the Université de Montréal. Natural language dialog, and the analysis and generation of text pose the current research challenges. Multilingual, natural language database queries are the topic of the MULTIQUEST project. A database interrogation tool which permits these queries will be designed and developed. The French firm Cap Sogeti Innovation is CRIM's partner in this project, which is also supported by the Franco-Québec cooperation program. A postdoctoral researcher from the Université Paris VI, Henri Béringer, is working at CRIM under the supervision of Drs. Kittredge of Université de Montréal and ORA, Boyer and Lapalme of Université de Montréal, and Sadri of Concordia University.

Engineering

The field of engineering likes to be at the front of computer technology and its applications. There is no lack of problems to solve and CRIM directs its efforts along with those of its members in this field.

An intelligent controller for the production of aluminum is the subject of a project where we examine the possibility of integrating AI techniques and classical control methods. The original ideas of this joint project with Alcan are knowledge modeling, the representation

Exploiting our Resources

The *Computer Research Institute of Montreal* (**CRIM**) is a non-profit organization whose members come from both industry and the academic community.

The people at **CRIM** are committed to the advancement of research in computer science. Whatever their rank or position, all of them contribute on a daily basis to ensure that their work promotes the development of goods and services.

In this respect, **CRIM** is somewhat like an agent for a precious resource which is continually fostered and renewed by scientists and specialists (research associates from industry), as well as by masters and doctoral students who stand to benefit from working with such outstanding research experts.

That precious resource is you—senior researchers, young Ph.D.s and graduate students—who are determined to help us meet the challenge of exchanging knowledge and achieving excellence.

For information, please contact:
Dr. Jacqueline Bourdeau
Scientific Administration
Electronic mail: bourdeau@bond.crim.cdn

Mailing address:
1550, de Maisonneuve Blvd. West
Suite 1000
Montreal, Quebec H3G 1N2

Computer Research Institute of Montreal

CRIM

L'essor d'une richesse

Le Centre de recherche informatique de Montréal (**CRIM**) est une corporation sans but lucratif qui regroupe des membres des secteurs industriel et universitaire.

Au **CRIM** travaillent des personnes qui se dévouent à l'essor de la recherche. Quel que soit leur rang ou leur poste, elles contribuent par leurs activités quotidiennes à faire en sorte que la recherche en informatique favorise le développement des biens et services.

En ce sens, le **CRIM** est un peu le courtier d'une richesse qui s'approvisionne et se renouvelle sans cesse auprès des scientifiques et de ceux qui collaborent à leurs travaux (associés de recherche des milieux industriels) ou encore de ceux qui trouvent auprès d'eux (étudiants de maîtrise ou de doctorat) l'occasion d'acquérir une formation de haut niveau.

La richesse, c'est vous, chercheurs chevronnés, jeunes docteurs ou étudiants gradués, vous qui êtes bien décidés à relever avec nous le défi de l'excellence et du transfert de connaissances.

S'adresser à :
Dr Jacqueline Bourdeau
Direction scientifique
Courrier électronique:
bourdeau@bond.crim.cdn

Adresse postale:
1550, bd de Maisonneuve ouest
Bureau 1000
Montréal, Qc H3G 1N2

Centre de recherche informatique de Montréal

CRIM

of physical elements in mathematical terms, and the identification of the most appropriate representations.

CRIM will also take part in the project Castor Plus along with the École Polytechnique de Montréal and seven industrial partners. The project involves generic research into numerical methods for hydraulic installations. The following subjects will be investigated: databases related to engineering, digital cartography and knowledge-based systems in engineering design. This project is directed by Drs. Tinawi, Camarero, Marche, and Soulié of École Polytechnique.

Systems Architecture for AI

Dataflow architectures for parallel applications in AI is the subject of another project in preparation with scientific consultants Dr. Gao of McGill University and Dr. Dennis of MIT. The goal of the project is to conceptualize a general-purpose dataflow architecture. One of the tools that will be developed for studying architectural concepts is a multiprocessor testbed simulator.

Laboratories

CRIM's two laboratories for Expert Systems and Software Engineering officially opened in April 1988. Headed by Mr. Charles Snow, the labs offer a favorable environment for the preparation and implementation of projects by providing a variety of hardware and software from which to choose. The labs also provide a forum for the exchange of ideas and the transfer of expertise and education. The labs offer short (one to three days) courses at the university level to members and non-members alike.

Equipment

CRIM's expert systems and software engineering laboratories make a variety of workstations and high-level software available to its employees and members. One can find Explorer II's from Texas Instruments, Sun 3/60's, a Sun 4/280, a MicroExplorer, a MicroVAX II and IBM microcomputers. Among the software available in these laboratories one can find the expert systems development shells ART, KEE, KnowledgeCraft, G2, Goldworks and NEXPERT Object. An IBM 9370, a VAX 8650 and several more MicroVAX II's are also available at CRIM.

Historical Background

CRIM is a non-profit corporation, financed principally by a support agreement with the Government of Québec and through the contributions of its members. Presently, there are five university members as well as 18 industrial members. The official inauguration of CRIM in October 1985 concluded the effort of a number of people working in universities, industries and the government to set up a computer science research centre where researchers from university and industry could collaborate. The development of CRIM has accelerated since January 1987, when it received the first installment of the financing promised by the Government of Québec. Today, there are about fifty employees at CRIM. CRIM initiates and develops research projects with members of its community and offers conferences and seminars, as well as access to its expert systems and software engineering laboratories.

CRIM has two principal objectives. First, in collaboration with industry, it identifies, promotes, and carries out computer research which will contribute to the development of goods and services in Québec. Second, it participates in the development of highly qualified personnel, especially at the doctoral level. The research domain which emerged as the most desirable in a survey of CRIM's members was artificial intelligence, specifically expert systems, as well as software engineering, man-machine interfaces, and computer communications.

Research Personnel

Dr. Renato De Mori is Vice-President and Director, Scientific Research, of CRIM. Dr. Jacqueline Bourdeau is Assistant Director, Scientific Research. In addition, CRIM has obtained the part-time collaboration of university professors, such as Dr. Bochmann of the Université de Montréal for projects in computer communications. There are a dozen holders of Master's degrees in computer science among the regular AI research staff at CRIM.

La recherche en intelligence artificielle au CRIM

par Diane Goupil et Jennifer Muise

Au CRIM (Centre de recherche informatique de Montréal), l'IA constitue une priorité et cet article présente le contexte, les grandes options de recherche ainsi que les projets qui s'y rattachent. Si l'IA constitue un thème central, les projets conduits au CRIM sont souvent reliés à plusieurs domaines, puisque le CRIM s'intéresse également au génie logiciel, aux interfaces personne-machine, à la télématique, à la CAO/FAO et aux circuits intégrés à très grande échelle. Ses laboratoires offrent un environnement favorable à la préparation ainsi qu'à la réalisation des projets.

Diane Goupil et Jennifer Muise sont des associées de recherche au Centre de recherche informatique de Montréal.

Les systèmes experts

Les systèmes experts et les systèmes à base de connaissances représentent une avenue de recherche-développement majeure au CRIM. On explore les différents modes de raisonnement, les méthodes et les outils de génie logiciel reliés aux systèmes experts, ainsi que les problèmes spécifiques à certains types d'applications.

Les problèmes de diagnostic sont étudiés dans le cadre d'un projet PRECARN avec Hydro-Québec et Alcan. Une "famille" de systèmes experts devrait couvrir le diagnostic d'un ensemble complet d'appareillage électrique, du point de vue de l'entretien ainsi que du

dépannage. Le projet porte sur un modèle d'acquisition de la connaissance pour la conception d'un système expert général de diagnostic de différents types d'appareils électriques. Ce projet est sous la direction scientifique des professeurs De Mori, Galiana et Lowther de l'Université McGill.

La modélisation de données a été étudiée avec la firme CGI dans le but de développer un système expert. Un prototype intelligent a été construit sous la direction du docteur Vaucher de l'Université de Montréal, afin d'identifier le potentiel de cette technologie pour les concepteurs de systèmes.

Avec la participation de PAPRICAN, le Centre canadien de recherche en pâtes et papiers, le CRIM recherche des solutions informatiques aux problèmes dûs aux dépôts de poix dans les usines de pâtes et papiers. Sous la direction du doctuer Allen, de PAPRICAN, l'équipe développe un système expert capable d'apporter des solutions à ces problèmes.

Comment un système expert peut-il produire et expliquer des prévisions météorologiques? Ce projet destiné aux aéroports canadiens est en cours avec Météoglobe Canada. Un prototype capable de prévoir la couverture et la hauteur des nuages bas a été réalisé. Le système visé intègrera la reconnaissance automatique de photographies par satellites ainsi que l'apprentissage automatique par l'expérience. Ce projet est dirigé par les professeurs Newborn de l'Université McGill et Zwack de l'Université du Québec à Montréal.

Le développement d'un système d'assistance pour la mise au point de modèles de simulation de vol fait l'objet d'une étude en collaboration avec CAE Electronique. Le but principal est de concevoir un système à base de connaissances qui doit assister les ingénieurs lors de la vérification et de la mise au point des modèles aérodynamiques des avions à simuler. Les professeurs De Mori et Bélanger de l'Université McGill assurent la direction scientifique du projet.

La reconnaissance de la parole

Bien des industries de pointe manifestent un intérêt tout spécial pour le potentiel de la reconnaissance automatique de la parole et de ses applications. Le docteur De Mori a développé ce domaine au CRIM à partir de plusieurs problèmes, notamment celui de la reconnaissance de la parole rapide dans le contrôle aérien et celui d'une interface vocale pour système expert. Le système présentera des questions à choix multiple et l'usager pourra donner sa réponse oralement plutôt qu'en utilisant le clavier. Plusieurs approches sont explorées, dont celle des réseaux neuronaux. Les projets sont réalisés en étroite collaboration avec deux partenaires industriels, CAE Electronique et IBM.

La robotique

Le CRIM s'est associé à l'Institut de recherche d'Hydro-Québec (IREQ) et au McGill Research Centre for Intelligent Machines (McRCIM) pour la recherche en robotique. Il participe à un projet PRECARN en collaboration avec ces deux établissements. Il s'agit du développement d'un robot intelligent guidé à distance par un opérateur afin d'automatiser l'entretien des lignes de tension. Le robot aurait une base de connaissances propre qui lui permettrait d'accomplir certaines tâches automatiquement, et de conseiller l'opérateur dans ses décisions. Ce projet est sous la direction des professeurs Daneshmend, Hayward, Levine et Zucker de l'Université McGill.

Le traitement des langues naturelles

Montréal possède une riche communauté de chercheurs en traitement des langues naturelles depuis les années 70 avec le grand projet de traduction automatique TAUM à l'Université de Montréal. Dialogue en langue naturelle, analyse et génération de texte sont des thèmes où de grands défis restent à relever. Au CRIM, le projet MULTIQUEST vise à concevoir et à développer un outil permettant une interrogation en langue naturelle et multilingue de bases de données. La firme française Cap Sogeti Innovation est le partenaire du CRIM dans ce projet qui reçoit le support du programme de coopération franco-québécoise. Un chercheur postdoctoral de l'Université Paris VI, Henri Béringer, travaille au CRIM dans ce projet réalisé sous la supervision des professeurs Kittredge de l'Université de Montréal et de ORA, Boyer de l'Université de Montréal, et Sadri de l'Université Concordia.

La résolution de problèmes en ingénierie

Le monde de l'ingénierie aime se tenir à la fine pointe de la technologie informatique et de ses applications. Les problèmes à résoudre ne manquent pas et le CRIM allie ses efforts à ceux de ses membres dans ce sens. Dans le cadre d'un projet de contrôleur intelligent pour le contrôle des opérations dans la production de l'aluminium, on examine les possibilités d'intégration des techniques de l'IA et des méthodes du contrôle classique. L'idée originale de ce projet, en collaboration avec Alcan est la modélisation des connaissances, la représentation des éléments physiques en termes mathématiques et l'identification des modes de représentation informatique les plus appropriés. La recherche effectuée dans le projet Castor Plus comporte un volet sur la représentation des connaissances dans le domaine des aménagements hydrauliques. Ce projet, qui se déroule en collaboration avec l'École Polytechnique et sept partenaires industriels, est sous la direction des professeurs Tinawi, Camarero, Marche et Soulié de l'École Polytechnique.

L'architecture d'ordinateurs pour l'IA

L'architecture de multiprocesseurs à flux de données pour les applications parallèles en IA consitue le thème de recherche du projet dirigé par les professeurs Gao de l'Université McGill et Dennis du MIT. Le but du projet est de conceptualiser une architecture de flux de données. Un des outils développés pour étudier les concepts architecturaux est un simulateur pour les systèmes à multiprocesseurs.

Laboratoires

En avril 1988, le CRIM inaugurait ses laboratoires de Génie logiciel et de Systèmes experts. Sous la direction de M. Charles Snow, les laboratoires offrent un environnement favorable à la préparation ainsi qu'à la réalisation de projets. Ils se veulent un environnement privilégié pour l'échange d'idées, le transfert d'expertise et la formation.

Équipement

Les laboratoires de Systèmes experts et de Génie logiciel du CRIM mettent à la disposition de ses employés et de ses membres plusieurs postes de travail et logiciels de haut niveau. On y retrouve des Explorer II de Texas Instruments, plusieurs Sun 3/60, un Sun 4/280, un MicroExplorer, un MicroVAX II, et des micro-ordinateurs

IBM. Parmi les logiciels disponibles dans ces laboratoires, on retrouve les outils de développement de systèmes experts ART, KEE, Knowledge Craft, G2, GoldWorks et NEXPERT Object. Un IBM 9370, un VAX 8650 et plusieurs autres MicroVAX II sont aussi disponibles au CRIM.

Contexte et historique du CRIM

Le Centre de recherche informatique de Montréal (CRIM) est une corporation sans but lucratif, financée principalement par une convention de subvention avec le gouvernement du Québec et par les cotisations des membres. Il compte présentement cinq membres universitaires ainsi que 18 membres industriels.

L'inauguration officielle du CRIM en octobre 1985 marquait l'aboutissement de l'effort de plusieurs personnes oeuvrant dans les universités, le gouvernement et l'industrie, afin de créer un centre de recherche en informatique dans lequel collaboreraient des chercheurs venant de milieux universitaires et industriels. Le développement du CRIM s'est accéléré depuis janvier 1987, date à laquelle il recevait le premier versement du financement prévu dans la convention de subvention avec le gouvernement québécois. Il compte aujourd'hui une cinquantaine d'employés. Le CRIM développe des projets de recherche avec les membres de sa communauté et offre des conférences, des séminaires, ainsi que l'accès à ses laboratoires de Systèmes experts et de Génie logiciel.

Personnel de recherche

Le docteur De Mori assure les fonctions de vice-président et de directeur scientifique du CRIM. Le docteur Bourdeau est directeur scientifique adjoint. De plus, le CRIM s'assure la collaboration de professeurs universitaires à temps partiel notamment celle du docteur Bochmann de l'Université de Montréal pour les projets en télématique. Parmi le personnel de recherche régulier travaillant en IA, on compte une douzaine de détenteurs de maîtrises en informatique.

AI Adventures at Autometrics

by James J. Dukarm

Aventures en intelligence artificielle à Autometrics

RÉSUMÉ: AAL Autometrics Associates Ltd. fait de la recherche et du développement dans le domaine des systèmes experts en temps réel et des systèmes experts interactifs par fréquences vocales. Cela nous a permis de gérer quelques applications comprenant des problèmes compliqués d'interface et nous avons lancé récemment un système de développement à partir de notre expérience. Nous faisons également du développement de systèmes experts sous contrat au moyen de coquilles commerciales. La majorité des applications sont du domaine de la foresterie et de la gestion des systèmes de génération d'électricité.

AAL Autometrics Associates Ltd. is a consulting and R&D company specializing in industrial computing, especially data collection, expert systems and voice technology applications. Located in Victoria, British Columbia, the company has a full-time technical staff of seven consisting of one Ph.D. (mathematical logic); one M.Sc. and three B.Sc. (computer science); one B.Sc. (mathematics) nearing completion of an M.Sc. (computer science and education); and one electronics technologist. Autometrics has been in operation since April 1985 and is a member of the Canadian Society for Fifth-Generation Research.

Research

Expert System Tool Development

Over the last three years, Autometrics has learned to implement small, fast rule-based systems in C running in a variety of environments ranging from simple ROM-based microcontrollers to minicomputers. This activity led to the release in June 1988 of our VRP Expert System Developer's Kit, which consists of some PC-based knowledge base development utilities and a run-time library for embedding rule-based modules in C programs.

James Dukarm is the president of Autometrics. He has a Ph.D. in mathematical logic and has been involved in the field of robotics and expert systems for 8 years.

We have received IRAP-M support from the National Research Council of Canada for one full-time and two part-time research positions to carry on further development of VRP. The main activities in that project are to :

- Enhance the expressive power of the rule base language without compromising its simplicity or significantly degrading the size and speed of VRP-based modules
- Improve the development utilities and add new ones for knowledge acquisition and rapid rule entry
- Port the development tools and run-time library to Unix, VMS and other strategically important environments.

Voice Data Collection

Early in the company's history, when we were more hardware-oriented, we prototyped several instruments for field and lab data collection under contract to the B.C. Ministry of Forests. In the course of analyzing the applications, we determined that in many instances direct entry of source data into a handheld computer in the field should be done by speech recognition rather than by keypad or expensive custom instrument. This was confirmed by means of field trials. It has taken a long time and many generations of prototypes to get all the elements of a viable product together and financed, but we believe that we are now close to having a good

handheld voice data collection system ready for marketing.

Systems Development

We have developed or are developing several expert systems using a variety of expert system tools, including Personal Consultant Plus, VP Expert, ESE, ADS, and our own VRP. The following is a summary of the most interesting items.

Reforestation Job Aid

We are preparing to release a reforestation "job aid" for field testing. The job aid is a ruggedized handheld computer running a very complex expert system written with our VRP expert system tool. The software consists of eight C programs and twelve rule bases (about a thousand rules) which call one another. Its job is to assist forestry field workers with site classification, generation of survey parameters, collection and analysis of survey data, site treatment recommendations and species selection for replanting. The project is subsidized by the Science Council of British Columbia. Ecologists and silviculturists from the B.C. Ministry of Forests and the Forest Engineering Research Institute of Canada have served as domain experts for knowledge base development, and will field test the unit. Eventually, speech recognition will be added to the job aid for data collection.

The reforestation job aid has demonstrated that our VRP expert system tool can be used for systems an order of magnitude larger than those for which it was intended While building the knowledge base for the job aid, we gained valuable insights and experience which have contributed to the further development of VRP, and to the design of knowledge acquisition utilities.

Heritage Guide

In connection with a consulting contract at B.C. Systems Corporation, we embedded a VRP expert system in a video and touch-screen interactive "Heritage Guide" program to provide an intelligent front-end for a tourist information facility. The main challenge in this project was to retrofit the rule-based module transparently into a rather odd host software environment.

Other Embedded Applications

We have used VRP to implement a variety of small embedded applications:

- Printer diagnostics running out of 32 K EPROM on a 2-MHz Z80 with 8 K bytes of RAM
- A telephone wine advisor with speaker-independent "yes/no/why/unknown" input and digitized speech output
- An experimental rule-based module to control polling of a power instrumentation network. The expert system, embedded in the firmware of a switchgear monitoring product we developed for Siemens Energy and Automation Inc., juggled polling priorities and system status information in real time to decide what command packets to transmit to which devices
- We developed a fast phonetics-to-text program for a PC-AT which recognized pieces of words in a stream of phonetic data and used a little VRP expert system to attach suffixes to words at two hundred words per minute. The program was done under NRC IRAP-L sponsorship for a client.

Forest Inventory Planner

We are using Texas Instruments Personal Consultant Plus to prototype an expert system for planning forest inventory projects. The B.C. Ministry of Forests is partially subsidizing the work and is furnishing the "experts."

What Next?

Our main research and development effort is to develop tools and techniques for building fast, small embedded rule-based systems. We are also looking for an opportunity to become involved in neural network applications.

As the company prepares to start its fourth year, we are paying more attention to the "bottom line" and shifting our emphasis from product development to product sales and consulting. Most of this will involve software and expert system development in our customary market areas — forestry, other resource industries, and the electric power industry.

AAAI-88: The Seventh National Conference on Artificial Intelligence

by Patrick Fitzsimmons, Stephanie Miller and Scott Goodwin

AAAI-88: la septième conférence nationale sur l'intelligence artificielle

RÉSUMÉ: Les organisateurs de AAAI-88 déclarent que le but de la conférence est "de promouvoir la recherche du plus haut calibre en intelligence artificielle et de promouvoir l'échange scientifique entre les chercheurs et les praticiens de l'intelligence artificielle." Un total de cent quarante-huit articles, neuf conférenciers invités et deux panels couvrent onze sous-domaines différents de l'intelligence artificielle: l'intelligence artificielle et l'éducation (un article), le raisonnement automatique (38), les modèles cognitifs (7), le raisonnement de sens commun (20), la représentation des connaissances (27), l'apprentissage et l'acquisition des connaissances (22), les architectures de machine et les languages de programmation (10), la langue naturelle (9), la robotique (7), les interfaces usager (4) et la vision (3).

22 - 26 August 1988, St. Paul, Minnesota

Introduction

The Seventh National Conference on Artificial Intelligence was held from the 22nd to the 26th of August 1988 in St. Paul, Minnesota. The organizers state that the purpose of the conference is "to promote research of the highest caliber in artificial intelligence (AI) and to promote scientific interchange among AI researchers and practitioners."

The AAAI conferences have grown steadily since their inception in 1980. From only two parallel sessions in 1980, the conference has grown to include: five parallel tutorial sessions over two days; four parallel paper sessions over four days; 15 workshops; invited talks; panel discussions; a major exhibition; two conference receptions and several vendor receptions. The size of the conference proceedings* has also grown, from one small volume in 1980 to two large volumes this year.

For a rookie conference attendee, there were far too many decisions to make during the conference. Too many things were going on all at once. There were four parallel technical sessions, invited speakers and panels, as well as vendor presentations and exhibitor booths (including AI publishers, expert system shells, AI machines, Lisp implementations, even a magician!). Tutorials were held in the two days before the sessions began, and workshops were held prior to and during the conference as well. It was an experience to see (and even meet) some of the people who wrote the textbooks from which we've learned. There really was something for everyone — from the academic, to the businessman, to the casual observer merely interested in AI, to the party animal.

Patrick Fitzsimmons, an M.Sc. student at the University of Alberta, is applying Theorist, a nonmonotonic reasoning system, to turf grass disease diagnosis to test its practical feasibility on a real-world problem. Stephanie Miller, a former M.Sc. atudent at the U of A, studied the use of specialists to accelerate general resolution-based theorem provers. Scott Goodwin, a Ph.D. student at the U of A, is studying nonmonotonic reasoning with emphasis on its application to temporal reasoning.

With all of this going on, we could not hope to cover the entire conference. Instead, we will report on those events that we attended and we describe some of the events not covered in the proceedings.

Organization

A total of 148 papers, nine invited speakers, and two panels covered 11 different sub-areas of AI. There were 850 papers submitted this year (up from 734 last year), but only 148 accepted (down from 155 last year). A hot area this year appeared to be learning and knowledge acquisition, with submissions up by a factor of two. The expert systems field is no longer a separate sub-area, having progressed to the point where they are included in the areas to which they contribute. AI has grown too large to be adequately covered in just one conference; specialized conferences, such as the new "Principles of Knowledge Representation and Reasoning" conference, are beginning to appear.

This year's conference filled the entire St. Paul Civic Center with one of the parallel paper sessions held in the nearby Radisson Hotel. Over 7 000 people were involved in the conference this year. This year's co-chairmen were Reid Smith and Tom Mitchell. The program committee was up to 79 people this year, from 68 last year, with 42 of them new people. Thanks are owed to these people and the many other organizers for taking the effort to get involved and producing an exceptional conference.

Papers

AI and Education

There was only one paper in this section, by Farrell, who described a system which tracks student assumptions and illustrates violations of them to the student as a self-education aid.

Automated Reasoning

There were a total of 38 papers in this section,

*AAAI-88 Conference Proceedings are published by Morgan Kaufmann Publishers, Inc., and are distributed in Canada through John Wiley and Sons Canada Ltd. (ISBN 0-929280-00-8, US$75, US$56.25 for AAAI members).

sub-divided into eight different sub-areas, three of which were the following:

Theorem Proving

There were four papers in the theorem proving section which described methods (and some systems) which attempt to: (1) speed up the proof process for various logics (goal-directed strategies by Dershowitz and Sivakumar, use of specialists by Miller and Schubert), (2) make the proof process more natural and intuitive (Groeneboer and Delgrande), or (3) extend a proof method to handle a larger domain (Jackson and Reichgelt), maintaining completeness while doing so.

Uncertainty

The uncertainty section had only three papers although a number of other papers dealing with uncertainty in specific areas were in other sessions. Zarley, Hsia and Shafer described a system based on the Dempster-Shafer theory using Markov trees for belief propagation. Lasky and Lehner talked about a formal equivalence between the Shafer-Dempster belief theory and assumption-based truth maintenance (incorporating a probability calculus). McLeish gave an answer to a question involving solutions to the entailment problem.

Planning

We attended four of the 12 planning talks. Hogge's talk on *Prevention Techniques for a Temporal Planner* was concerned with a seldom-considered aspect of planning, viz., goal prevention. Dean and Boddy considered planning under varying time constraints. Pednault presented a methodology for solving planning problems involving actions whose effects are context dependent. Finally, Kaelbling's talk on *Goals as Parallel Program Specifications* focused on a formalism based on situated automata theory that addresses real-time planning in dynamic worlds.

The other five sub-areas were: automatic programming (four), constraints (two), rule-based reasoning (two), search (seven) and truth maintenance systems (four).

Cognitive Modeling

There were seven papers in this section (of which we attended three presentations) but no real "theme" existed among them. Lytinen and Moon described a system which learns the grammar rules of a second language (German) from English instructions. Reeves talked about how belief conflict patterns (and character assessments) can be used in narrative understanding. Fisher described a memory model that accounts for human preference for particular levels of generality and typicality effects.

Commonsense Reasoning

This section had 20 papers divided into three sub-areas: qualitative reasoning (12), design (four) and diagnosis (four). One of the two best papers, by Williams, was in the qualitative reasoning sub-area (see the Best Paper Awards section).

Knowledge Representation

For this topic, there were 27 papers altogether, divided into five sub-areas: inheritance (six), nonmonotonic logic (seven), nonstandard logics (four), temporal reasoning (six) and miscellaneous (four). Below we highlight the talks that we attended.

Inheritance

MacGregor described LOOM, a logic-based deductive pattern matcher for a knowledge representation system. Padgham discussed how to model and represent type information for use in reasoning with defaults by separating the definition of a type into a core and a default (optional) characteristic. Cohen and Loiselle discussed plausible inference rules in semantic networks based on generalizing ISA inference structures to other structures. Haugh presented a very elegant tractable method for representing and reasoning about multiple defeasible inheritance. Horty and Thomason's paper considered heterogeneous inheritance networks. Their Pennsylvania Dutch example stirred much controversy. It is clear that there still is a *clash of intuitions*. A deeper understanding of inheritance than is afforded by syntactic graph traversers is necessary before real progress can be made. Equally perplexing was Touretzky's and Thomason's *Nonmonotonic Inheritance and Generic Reflexives*. It's amazing how many special characters must be used to specify a language for saying "elephants love themselves".

Nonmonotonic Logic

Konolige presented a method of nonmonotonic reasoning based on inferencing from specific evidence. Przymusinski looked at the relationship between nonmonotonic reasoning and logic programming and tied them together based on a new declarative semantics of logic programs that has been shown equivalent to four major nonmonotonic formalisms.

Nonstandard Logics

Dalal's talk focused on revising knowledge in the face of new, contradictory knowledge. Bonner's talk on *A Logic for Hypothetical Reasoning* discussed a logic for reasoning with "hypothetical rules" such as "You are eligible for citizenship if your father would be eligible if he were still alive." Patel-Schneider describes a four-valued semantics for a tractable logic that supports subsumption.

Temporal Reasoning

All but one paper in this section, by Ladkin, were on reasoning about change, persistence or causal reasoning; as opposed to pure temporal reasoning. The Yale Shooting problem was also prominent in many papers. Ladkin talked about a first order theory for detecting inconsistencies in arbitrarily quantified formulas in Allen's interval calculus. Myers and Smith tackled the problem of persistence in derived formulas, as most work has been done on non-derived formulas. Hanks designed a representation and a calculus for determining an agent's belief in a proposition at a point in time using the tendency of the proposition to persist over time. Morris attempted to fix a problem with the single extension solution to the Yale Shooting problem by proposing a new belief formalization which can handle new conflicting information. Morgenstern and Stein gave a theory of generalized temporal reasoning to handle temporal projection and explanation, which can be

applied to the Yale Shooting problem. Dean and Kanazawa described a theory of causal reasoning under uncertainty to handle persistence.

Learning and Knowledge Acquisition

There were 22 papers in this section, divided into five sub-areas: learning (six), analytic learning (six), empirical learning (five), formal results in learning (four) and knowledge acquisition (one). One of the two best papers, by Minton, was in the analytic learning sub-area (see the Best Paper Awards section).

Machine Architectures and Computer Languages

This section had a total of ten papers divided into three sub-areas: architectures and languages for problem solving (three), production systems (four) and miscellaneous (three).

Natural Language

There were nine papers in this section, two in syntax, three in speech and four in dialogue.

Dialogue

The papers here dealt with several areas in dialogue: incomplete lexical information — how to understand the text in spite of it and how to incorporate the words for future use (Jacobs and Zernik); interpreting temporal relations from narratives, including the time relations of tenses, and relations between clauses and sentences (Song and Cohen); handling semantic ambiguity by acquiring missing information from subsequent input (Moerdler and McKeown); and using the user's knowledge to determine the content of an answer (Chin).

Robotics

This section had only seven papers, four in robotics and three in integrated robotics. Buckley showed how, in the interface between a compliant motion programming system and a compliant motion control system, a robot can be modeled as a damping spring. Stansfield described reasoning about grasping during the precontact stage using vision and touch as an aid to selecting hand preshapes and reach parameters. Kuipers and Byun described a robust, qualitative method of spatial learning based on a successive layered model of procedural knowledge, a topological model and metrical information.

User Interfaces

This section had only four papers. We were unable to attend any of these.

Vision

This section had only three papers. Last year there were 14 papers in this section. We were unable to attend any of these.

Invited Talks and Panels

Terrence Sejnowski reminded us of the size of the task of machine intelligence in his very interesting talk *What is the Computational Power of the Brain?*

Takeo Kanade convinced us of the need for going back to the basics in computer vision, and took us there, in his autonomous robots talk *Computer Vision: From Ad Hocs to Science.*

John McDermott expressed his views on tools for knowledge transfer in his talk *Exploiting Task Structure to Automate Knowledge Acquisition.*

Earl Sacerdoti spoke about the more practical side of AI and its financial successes in his talk *In Search of the Real World: Practical Utility of AI.* He made some predictions of the future, noting that investment has paid off, and that we can expect to see investment shifting towards hardware.

Michael Brodie surprised his audience when he finished his talk, *Future Intelligent Information Systems: AI and Database Technologies Working Together*, early and went on to discuss the moral responsibilities of AI researchers. He advised the audience to consider the ultimate uses of their research and to think about the kind of world they are shaping for their children.

Harold Cohen treated us to a refreshingly different use of AI — art — in his talk *How to Draw Three People in a Botanical Garden*. He had some wonderful pictures to show us, some of which are included in the proceedings, but, alas, not the coloured ones.

We were unfortunately unable to attend three of the invited talks: Andrew Witkin's *Physically-Based Modeling for Vision and Graphics*, David Gossard's *Challenges for AI in Future CAD/CAM Systems* and Carver Mead's *Analog VLSI Models of Neural Systems.*

In the panel led by Edward Feigenbaum, *Results from a Study of the First Wave of Expert System Applications to Business*, we found out about some of the approximately 2000 AI expert systems that are actually in use in the "real world" (1500 in the U.S., 250 in Europe and 250 in Japan), including what happens when American Express accepts (or rejects) a charge you try to make. Astounding figures were presented on the money that expert systems have saved for companies. For instance, DEC has saved $70-100 million per year, IBM has saved $10 million per year, Dupont has saved $100 000 per year and American Express has saved $27 million per year. Expert systems have vastly increased productivity, and hence, our national growth. The only negative comment towards expert systems was that there are too many similar expert systems out there with different names.

In the panel led by Howard Shrobe, *Where We Are and Where We Are Going*, the history of AI and its future challenges were reviewed. The issue of the integration of AI sub-areas was discussed.

AAAI-88 Presidential Address

AAAI President, Dr. Raj Reddy of CMU, gave a resoundingly optimistic talk (entitled *Towards a Theory of Knowledge*) about the goals of AI, what we have accomplished and what we have left to do. He even suggested that the so-called "AI winter" might be nearing an end as more and more AI applications make their way into the business world and start showing a return on investment. He raised the issue of AI having to be accountable for itself since there are fewer people willing to blindly trust AI. He described several rather impressive AI applications (mostly from CMU) for chess, speech understanding, computer vision and expert systems, effectively interspersing slides and videos of the systems into his talk.

Next, Dr. Reddy presented what he called the *Principles of AI* that have come out of lessons learned from years of AI research:

(1) bounded rationality implies opportunistic search,
(2) a physical symbol system is necessary and

sufficient for intelligent action
(3) the magic number: 50,000 ± 20,000, (expert knowledge ± chunks of knowledge)
(4) search compensates for the lack of knowledge
(5) knowledge eliminates the need for search.

In conclusion, Dr. Reddy talked about the future challenges of AI and the social implications of AI. Some of the future accomplishments he mentioned to look forward to were: for AI chess programs the next step is to beat the world champion. Look toward some amazing mathematical result from a computer. How about an automatic translator telephone system that allows multi-language conversation (eg. Chinese at one end and French at the other end)? What about accident avoiding automobiles and intelligent cruise control? How about systems that can read text and then answer questions, in other words, learn by observation? Finally, what about self-replicating systems? Certainly, his talk was one of the highlights of the conference.

Best Paper Awards

Beginning in 1982, AAAI Best Paper Awards have been given in recognition of "papers that report important, substantial research in an exemplary manner." Of the 148 papers selected by the Program Committee for AAAI-88, 11 were nominated for the Best Paper Award. Of these, two were identified as clear winners.

Steven Minton's paper, *Qualitative Results Concerning the Utility of Explanation-Based Learning*, is concerned with the effectiveness of PRODIGY, "a learning system that searches for effective control knowledge". This paper reports on experiments measuring PRODIGY's performance.

Brian William's paper, *MINIMA: A Symbolic Approach to Qualitative Algebraic Reasoning,* discusses a "hybrid algebra, called Q1, [which] allows us to select abstractions intermediate between traditional qualitative and quantitative algebras". This algebra is the basis of the qualitative symbolic algebra system, MINIMA.

The selection of these papers reflects, in part, the desire of the Program Committee to encourage researchers to go beyond theory and implementation to perform and report experimental results.

The award winners received a cash prize and were invited to submit a revised version of their paper to the AI journal.

The other nine nominees were: Jonathan Amsterdam for *Some Philosophical Problems with Formal Learning Theory;* Brian Falkenhainer and Kenneth D. Forbus for *Setting Up Large-Scale Qualitative Models;* John J. Grefenstette for *Credit Assignment in Genetic Learning Systems;* Leslie Pack Kaelbling for *Goals as Parallel Program Specifications;* Benjamin J. Kuipers and Yung-Tai Byun for *A Robust, Qualitative Method for Robot Spatial Learning;* Mark F. Orelup, John R. Dixon, Paul R. Cohen and Melvin K. Simmons for *Dominic II: Meta-Level Control in Iterative Redesign;* David B. Searls for *Representing Genetic Information with Formal Grammars;* Reid G. Simmons for *A Theory of Debugging Plans and Interpretations;* L. G. Valiant for *Functionality in Neural Nets.*

Receptions

There were two conference receptions, each offering plenty to see, eat and drink (well, maybe not that much to drink).

For the first reception, everyone was bused out to the Minneapolis Zoo, about 20 miles south-east of the St. Paul Civic Center. While at the Zoo, everyone was free to tour and see all the wonderful wildlife.

The second reception was held at the historic Landmark Center, right across the street from the St. Paul Civic Center. Everyone was treated to live music and were free to tour the building.

In addition, many of the vendors sponsored private receptions with free food and drinks as well as the customary propaganda. The Sun Microsystems reception was held at the Science Museum of Minnesota, complete with free food, drinks, complete access to the museum exhibits and attendance at the Omnitheatre to view a showing of the IMAX film "Seasons". One of the museum exhibits was a room full of giant life-like robot dinosaurs that were quite scary after consuming a few drinks.

Two private exhibitions, limited to those who pre-registered for them, were at UNISYS and the 3M Center. For the 3M Center exhibit, people were bused out to the 3M Center's Complex to view five exhibits: KLUE (The Knowledge Legacy of the Unavailable Expert), a "do-it-yourself" diagnostic expert system tool for manufacturing that allows process experts to "lay out" their problem-solving strategies in the form of a decision graph which is then used by an Advisor program, tailored to the needs of the process expert; EE&PC, a product environmental risk assessments system that evaluates product formulations according to criteria determined by expert environmental engineers; 3M Product Catalog on a Hypercard stack; Video Disc Demonstrations of the capabilities of 3M's optical media, including prerecorded video and CD formats, re-writeable (650 megabyte capacity), 12 inch write-once (WORM, one gigabyte capacity per side) and erasable (500 megabyte capacity) formats; and CD-ROM demonstrations of the applications of commercial CD ROM discs (non-erasable, 552 megabyte capacity).

Review of Exhibits

Present at the exhibition this year were such companies as: Apollo, Apple, Data General, DEC, Hewlett-Packard, IBM, Sun Microsystems, and Texas Instruments, representing only about one third of those companies marketing AI technologies or applications.

The emphasis this year was on expert systems and hardware. Several expert system shells were on display.

Publishers such as Morgan Kaufmann, Addison-Wesley, Springer-Verlag, John Wiley & Sons, Prentice-Hall, etc. were in full force, offering a wide selection of AI publications (and some free samples!).

This year, credit card-like plastic cards with names and addresses came as part of each registration package. These cards fit into the identification badges and exhibitors could quickly take impressions of them using credit card style charging machines, thereby relieving the necessity of filling out address forms for exhibitors.

Workshops*

According to *Canadian AI Magazine* [1], there were 15 workshops. *Uncertainty in AI* was held August 19-21; *Neural Architectures for Computer Vision* was held on

*AAAI-88 Workshop Proceedings are available from AAAI (445 Burgess Drive, Menlo Park, CA 94025, U.S.) for US$20 plus $2.40 shipping and handling.

August 20; *Parallel Algorithms for Machine Intelligence and Pattern Recognition* was held August 20-21; *Integration of Knowledge Acquisition and Performance Systems* was held on August 21; *Explanation* was held on August 22; *AI and Hypertext: Issues and Directions,* and *Case-Based Reasoning* were held on August 23; *AI and Music, AI in Design,* and *Blackboard Systems* were held on August 24; *AI in Process Engineering, AI in Production Planning and Scheduling,* and *Automating Software Design: Current Directions* were held on August 25; *Databases in Large AI Systems* was held on August 26; *Conceptual Graphics* was held on August 27.

Commendations and Complaints

Busing was necessary between the conference and the various hotels as well as to and from the receptions and private exhibitions. Congratulations to all those who helped organize the extremely efficient bus schedules.

T-shirts, as usual, ran out within half a day. This year's design was the triangle and wave symbol from the AAAI-88 poster on black background, staff T-shirts were the same except on light grey background.

There was absolutely no seating available in the lobby or outside the session rooms at the St. Paul Civic Center which left people sitting on the floor if they needed to rest and decide which talks to see next. This was especially a problem upon arriving and receiving a registration package full of things to go over but nowhere to sit down.

Having the conference parallel sessions held at both the St. Paul Civic Center and the Radisson Hotel, about three blocks away, made it quite difficult to get to successive sessions if you had to travel between the two locations.

Reference

[1] Canadian Society for Computational Studies of Intelligence, *Canadian Artificial Intelligence,* No. 16, July 1988.

Book Reviews

edited by Graeme Hirst

Automated Reasoning:
33 Basic Research Problems

Wos, Larry
[Argonne National Laboratory]

Englewood Cliffs, NJ: Prentice Hall, 1988, xiii+319 pp
Paperback, ISBN 0-13-054522-X, US$12.00

Reviewed by
Howard Hamilton
Simon Fraser University

Howard: I still haven't reviewed that book for *Canadian AI*. I've had it about nine months.
Jim: What's the book?
H: It's called *Automated Reasoning: 33 Basic Research Problems* and it's by Larry Wos.
J: *[Tries a sip of coffee ... too hot]* What's the problem?
H: I read about a third of it and skimmed the rest, but I just can't seem to finish reading it. I guess part of the problem is that it's about theorem proving, which is not my bag.
J: So why did you agree to review it?
H: Well, I was misled by the blurb on the back cover. It said that the book presented 33 basic research problems in automated reasoning in a manner that the interested generalist could understand. As you may recall, about then I was in search of a Ph.D. thesis topic and the idea of someone listing 33 research problems sounded very appealing. I was interested in commonsense reasoning and I thought the book would have something for me — perhaps put the problem of automating commonsense reasoning in a context of automating all reasoning, or something like that. The preface kept up the illusion. Look here in the first paragraph: "Automated reasoning is concerned with computer programs that reason. The reasoning may be logical, probabilistic or commonsense" (p.1).
J: What went wrong?
H: After the preface and introduction, Wos forgets about probabilistic and commonsense reasoning. Apparently, Larry Wos is one of the best known researchers in the field of automatic theorem proving. He interprets "automated reasoning" to mean what I call "theorem proving". *[Leafs through the book.]* See, on page 65 he states that "the name *automated reasoning* was first introduced in 1980 ... to reflect the expanded use of those programs that were originally designed to prove theorems". Anyhow, the focus of the book is theorem proving and not "automated reasoning" in the broad sense that I think of it.
J: *[Takes a sip of coffee]* Pretty exciting stuff, eh?
H: Wos thinks so. On many occasions, he stops and does a bit of "cheerleading" for the field; here are two examples from page 2:

> "The beauty of a theorem from mathematics, the preciseness of an inference rule of logic, the intrigue of a puzzle, and the challenge of a game — all are present in the field of automated reasoning.
>
> "The various successes, numerous uses, and diverse applications of automated reasoning strongly suggest that the future is indeed most promising and exciting."

Reading such comments once or twice is okay, but by about the fifth time (p.11), I was getting pretty sick of them.
J: So what are the 33 problems?
H: The 33 problems seem a rather mixed bunch to someone from outside the field of theorem proving. I'll give you an example. One of his problems is "induction". Here, see, number 29:

> "What criteria should be used to select the property on which to base an induction argument?"

Well, it seems to me that induction is a whole field of study with lots of previous work in philosophy and AI, not just one research problem, although it might appear that way to someone preoccupied with building theorem provers. The problem is not described from the perspective of commonsense or probabilistic reasoning; it is simply considered with reference to proving theorems. But problem number 13 is at the other extreme:

> "What criteria should be used to decide when to permit and when to avoid demodulation during application of the inference rules of hyperresolution, UR-resolution and hyperparamodulation?"

That sounds like one suggestion for research based on a specific paper he may have written. Indeed, when I checked through the references, I found that the term "demodulation" was introduced by Wos *et al* in 1967 and "hyperparamodulation" (good thing I don't have to spell that word very often!) was introduced by Wos *et al* in a 1980 paper.
J: *[Takes another sip of coffee]* What are the rest of the problems like?
H: Most of them are like number 13 — you know, only theorem-provers need apply. The exceptions sound like whole fields of research to me. For example, number 33 is "reasoning by analogy":

> "Bearing in mind that reasoning by analogy is only one of many types of reasoning that people use to answer questions, what criteria can be applied to a question to cause an automated reasoning program to choose wisely to employ this type of reasoning?"

But most of the rest are about theorem proving. Here's number 15:

> "What is the appropriate theory of demodulating across argument and across literal boundaries — a theory similar to the current use of demodulation or similar to that for complete sets of reductions — to replace certain predicates by other predicates and certain collections of literals by other collections?"

I felt that saying the problems were accessible to the generalist was misleading. About 28 of them are about theorem proving and the other 5 cover some other fields of research. So, the book is really for grad students looking for research topics in theorem proving.
J: But, given that objective, is it a good book?
H: The book should have been a lot shorter. There just isn't enough content to justify 319 pages. Look at how padded it is. Chapter 3: a simple list of the 33 problems, interspersed with about a paragraph of commentary

each. Chapter 5: the same 33 problems again, but with a page or two of commentary each.

J: Is the rest good stuff?

H: Nah, it's incredibly repetitious. I found it hard to make progress through the book because I kept saying, "Oh I've read this part before." And invariably I was right. I had read it, almost word for word, in a previous chapter. Many statements are made three or more times. The book starts with a preface. Almost everything from the preface is repeated in chapter 1, the introduction; then almost everything in chapter 1 is repeated somewhere in the contents of chapters 2 through 6. Whoops! Chapter 3 isn't a content chapter because everything in it is repeated in chapter 5. Okay, so the content chapters are 2, 4, 5, and 6. Finally, chapter 7, the so-called conclusion, simply repeats the stuff in chapter 1. Now, that's repetition!

J: All books have some repetition; how bad is the repetition here?

H: I tell you — too much of this book comes from the copy command of a word processor. I'll give you an example dealing with the uses of automated reasoning. Going once, page 2, chapter 1:

> "in answering open questions from mathematics [Winker79, Winker81, Wos85b] and from logic [Wos83b, Wos84b], designing superior logic circuits [Wojciechowski83], validating the design of existing circuits [Wojcik83, Hunt85, Hunt86], and verifying the correctness of computer programs [Boyer84b, Boyer84a, Good82]."

Going twice, page 5, chapter 1:

> "in answering open questions from mathematics [Winker79, Winker81, Wos85b] and from logic [Wos83b, Wos84b], designing [Wojciechowski83] and validating [Wojcik83, Hunt85, Hunt86], logic circuits, and proving claims made for computer programs [Boyer84a, Boyer84b, Good82], for example."

Going three times, page 15, chapter 2:

> "in answering open questions from mathematics [Winker79, Winker81, Wos85b] and from logic [Wos83b, Wos84b], designing superior logic circuits [Wojciechowski83], validating the design of existing circuits [Wojcik83, Hunt85, Hunt86], and verifying the correctness of computer programs [Boyer84a, Boyer84b, Boyer84c, Good82]."

(Sold? Nope, just bored). And we still haven't heard the end of it — see page 106 (chapter 4) and page 250 (chapter 7). You can imagine that such repetition can lead to errors and inconsistencies. Look at the last list of references in the three example passages and you will see that it's hard to keep everything consistent. I think we learned something about this in database class! The worst thing is that this is not an isolated example. Many other phrases and sentences are repeated over and over again, as well.

J: So, what's good about the book?

H: Wos is extremely knowledgeable about his field. Some of the introductory material in chapter 2 is good and there is some interesting historical stuff about how Wos invented the "set of support" strategy for resolution theorem proving in 1964. Chapter 6 gives a set of 17 test cases for theorem provers and relates them to the "basic research problems". Also, the bibliography might be of use to the reader. And the index is great — it's 26 pages (fine print) of well-organized detail that enabled me to locate any point I remembered from reading the book. However, overall, I'd say you'd be better off buying Wos' text on the subject instead. (As Wos suggests on p.xii, "A full treatment of automated reasoning is given in the companion to this book, *Automated Reasoning: Introduction and Applications*, written by Wos, Overbeek, Lusk, and Boyle, and published by Prentice-Hall"; the same book is also recommended on p.3, p.8, p.17, . . .) If you like the text and its perspective on the field of theorem proving and want a list of research problems, then buy *33 Basic Research Problems*.

J: *[Stares at the bottom of an empty Styrofoam coffee cup]* What about the format, with its list of research problems?

H: The book is an experiment, and on balance I'd say it is a failure. I suspect that it will not make theorem proving appreciably more accessible to generalists. Originally, I thought the list of questions was a great idea — just what I needed to find my topic. Wos even suggests that "the difficulty of obtaining a complete solution to one of the problems is probably equal to that of finishing a Ph.D. thesis, and, in fact, the problems are suggested as thesis topics." But when I see "reasoning by analogy" as one research problem (number 33) and I know how much research is actually being done on the topic, I realize there's no high road to scientific discovery. One researcher has a limited perspective and can describe only the open problems relevant to the research he knows well. Given the amount of time and effort required to produce such a book (and the attendant delay since the research problem was first observed), I don't see much use in writing any more "list of problems" books.

J: *[Tosses coffee cup at garbage can; misses]* Why don't you write a review just saying what you've told me?

H: Nah, I don't think anyone would want to read a review like that. I'll just send the book back and let someone with more interest in theorem proving review it.

Howard J. Hamilton is a Ph.D. candidate in the School of Computing Science at Simon Fraser University. He is currently working on "Baconian induction of qualitative models" under the supervision of Nick Cercone. The role of "Jim" was played by Jim Delgrande, who is a faculty member at SFU.

Applications of Expert Systems: Based on the Proceedings of the Second Australian Conference

Quinlan, J. Ross (editor)
[New South Wales Institute of Technology]

Glasgow and Sydney: The Turing Institute Press
in association with Addison-Wesley, 1987, xvi+223 pp
Hardbound, ISBN 0-201-17449-9, US$24.95

Reviewed by
Dick Peacocke
Bell-Northern Research

In a collection of papers on expert systems you might expect to see multiple versions of a basic introduction to how expert systems work. You do not find that in this book. The book contains a tightly-packed collection of papers from the Second Australian Conference on Applications of Expert Systems held in May 1986. It stresses applications of expert systems through discussion of case studies and by considering the characteristics of particular types of application domain. The book does not contain new research results or detailed designs of the systems in the case studies. Its contribution is in showing some of the form, structure,

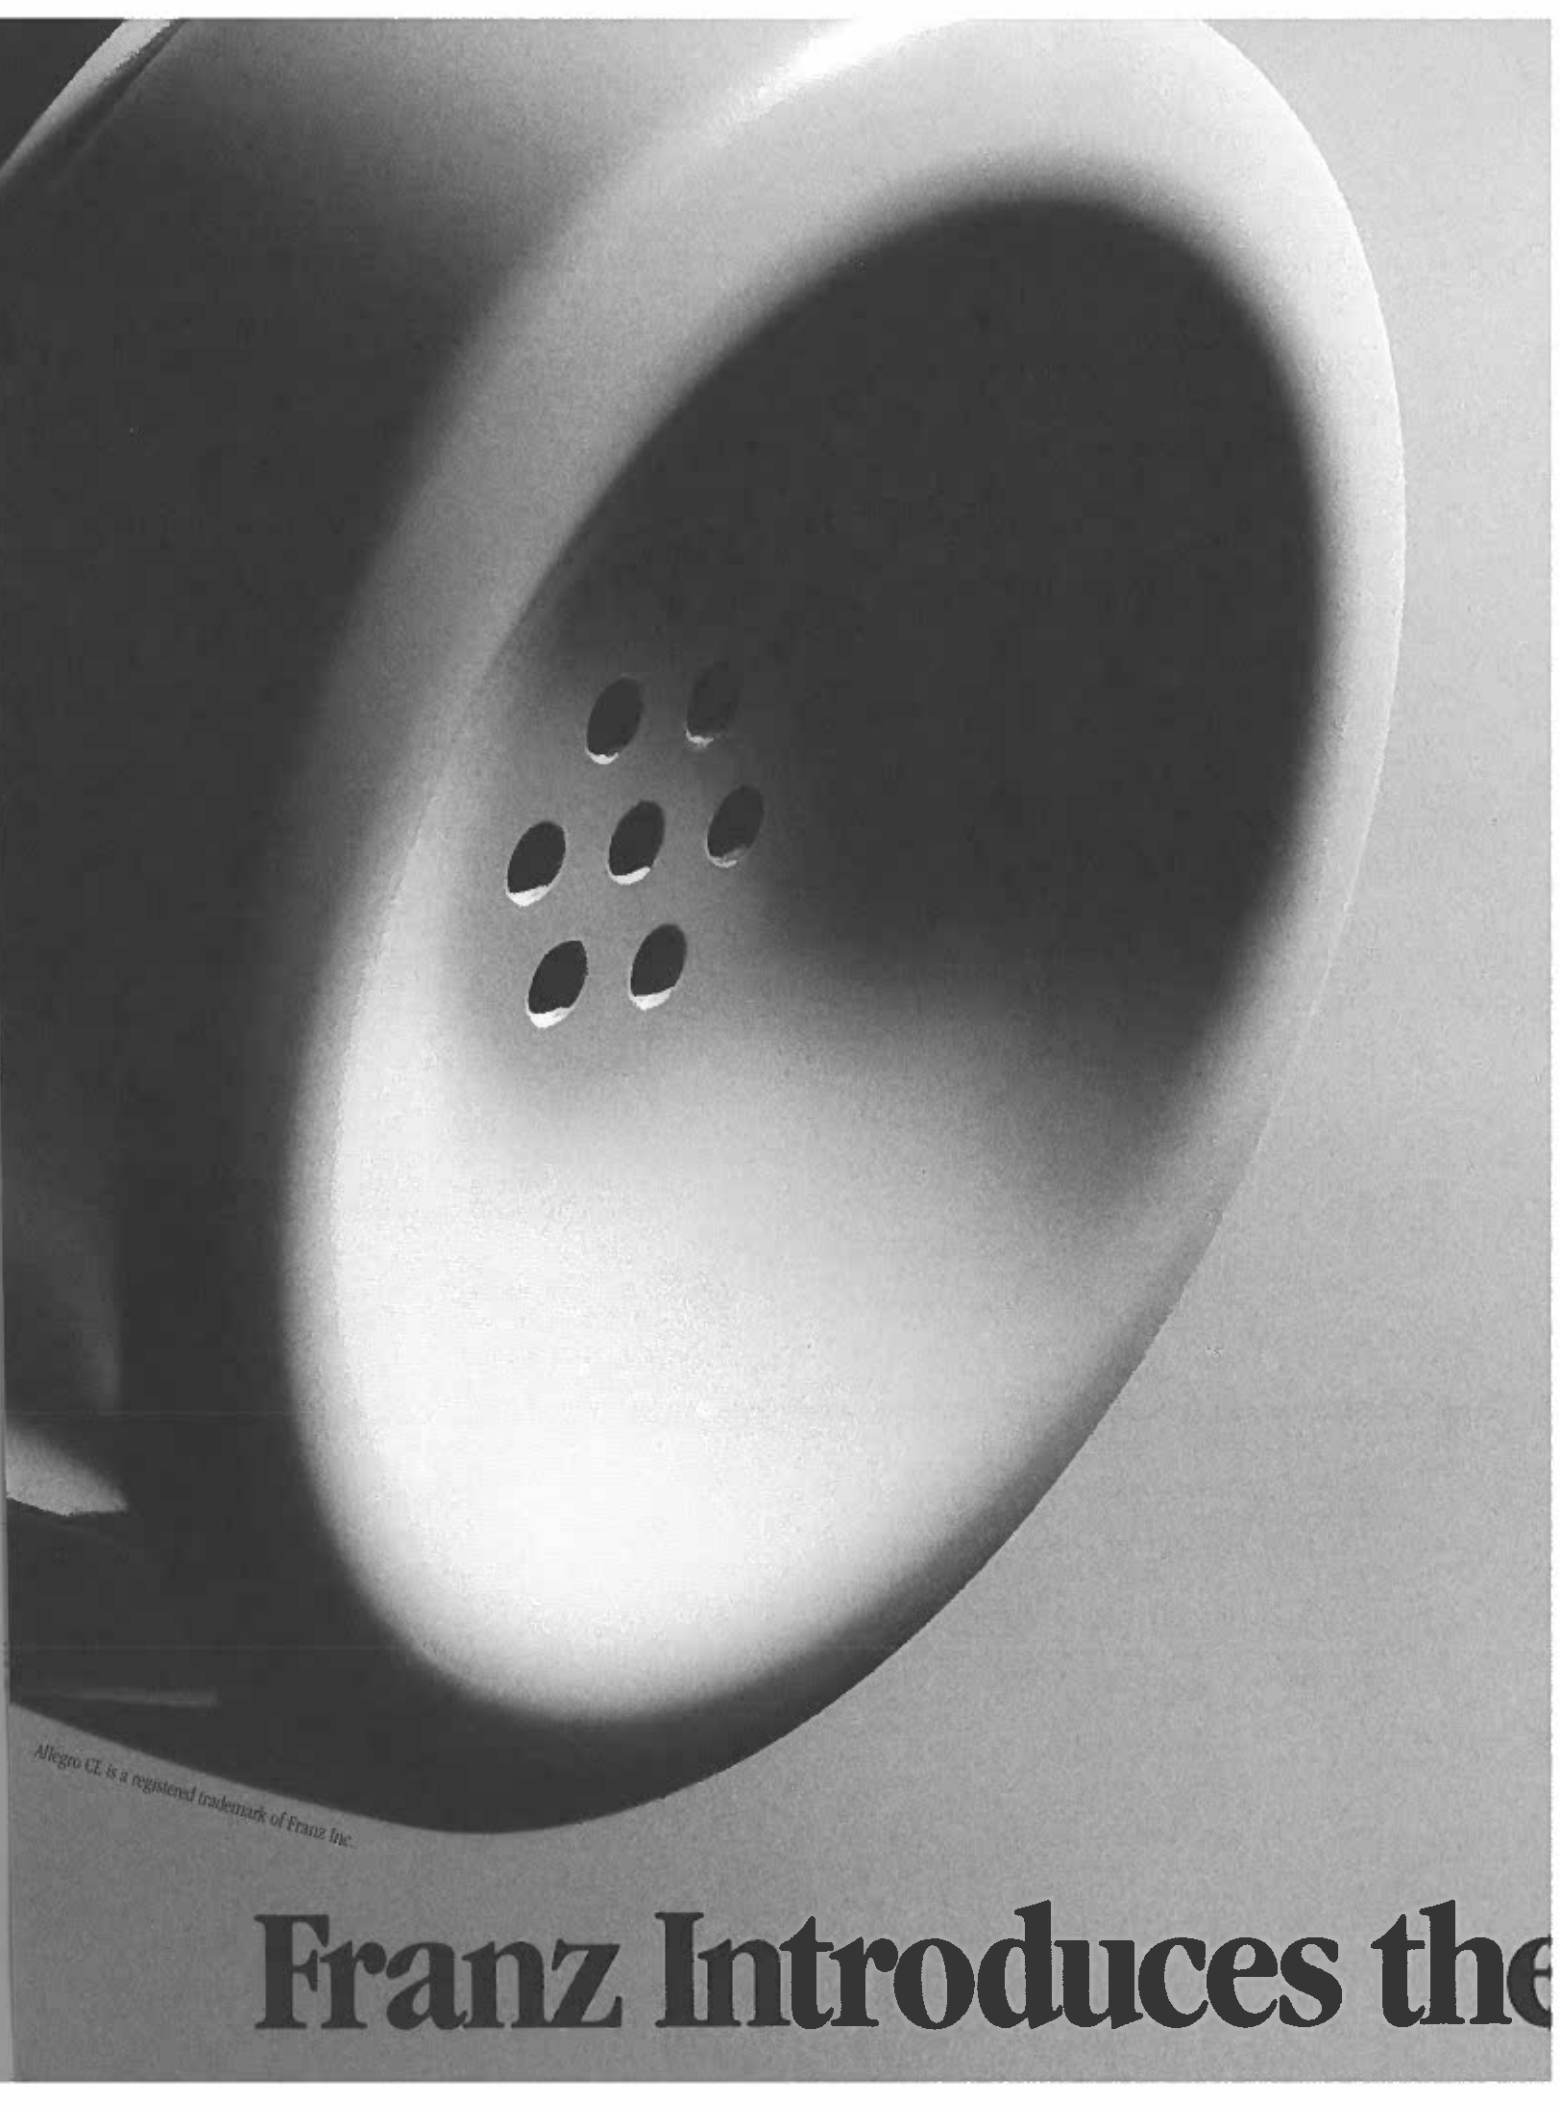

Allegro CL is a registered trademark of Franz Inc.

Franz Introduces the

Franz Inc. · 1995 University Avenue · Berkeley, CA 94704 · FAX: (415) 548-8253

Sure, we've got the most portable Common LISP in the business—micros to mainframes. And our Allegro CL® is setting new records for speed, compactness, and reliability. But when you're developing serious commercial applications, you need more.

"More than any other supplier, Franz has been incredibly responsive to all our needs, in record time."

Jeff Fox, VP Engineering—Silc Technologies, Inc.

If you have a LISP problem or need an enhancement, you have a choice: you can spend precious development time trying to get an answer from a hardware company, or you can come to the source.

"By far, the best customer support of any software vendor I've ever dealt with. Their open attitude and attention to our needs has been amazing."

Gordon Kotik, President—Reasoning Systems, Inc.

When you call customer service at Franz, you get a LISP expert. If you need a special feature or enhancement, you'll get it, fast—by dealing direct with the person who actually developed that part of Allegro CL.

Experience the ultimate in LISP support. Pick up your own Franz accessory, and call: (415) 548-3600.

FRANZ INC.

Ultimate LISP Accessory.

and limitations of the knowledge in various expert systems. It does that well.

The book is divided into three parts, and contains a total of 12 papers. Part One, Artificial Intelligence and Expert Systems, contains a single paper by Patrick Winston. It places the expert systems field within the broader context of artificial intelligence. Rule-based expert systems are introduced using the grocery bagging example, and the basic ideas behind XCON are explained. The paper contains advice for those new to the field on how to select a good application. The advice is good, but does fail to emphasize an essential point: the application must be needed and useful. Winston also summarizes different perspectives on tool evaluation; the bottom line is "Does the expert system solve the problem?"

Part Two, Expert Systems: Domains and Tools, contains six papers addressing legal, architecture and building applications plus real-time application requirements, and interfaces between expert systems and databases. The law is a difficult domain for expert systems because it relies on interpretation, argument, and weighing of conflicting factors. An excellent paper by the late Donald Waterman and two other authors presents systems for assessing product liability claims and for asbestos litigation. The next three papers address a variety of applications in the area of architecture and building. There is a paper which describes a system for analyzing the ability of window frame designs to exclude rain. It reasons about shapes and positional relationships, making extensive use of graphics. The next paper contains examples of expert systems for checking a design against a set of building regulations, for designing reinforced concrete retaining walls, and for spatial layout.

The following paper also focuses on establishing compliance with regulations, this time in the form of a fire safety code. A very different class of domain, real-time applications, is described in the next paper. A number of special requirements are identified for systems that must promptly assess circumstances in the light of changing information from the environment. The discussion is useful, but it never quite comes to grips with the problem in hand, identification of radar emitters. (This radar identification problem is encountered again in Part Three.) The last paper in Part Two presents a generic shell for applications that must access large amounts of data in an existing database. The shell has been used to implement an expert system for vehicle fleet maintenance.

Part Three, Knowledge Acquisition and Management, concentrates on inductive knowledge acquisition techniques. Expert systems often need large amounts of knowledge in order to achieve a satisfactory level of performance. In the first paper of the section, Donald Michie discusses the learning by examples approach to knowledge acquisition, based on inductive inference. He traces the approach from the pioneering work done by Hunt et al* and carefully follows various threads through to Quinlan's ID3 algorithm. The history is interesting and well written. Then the paper merges into a commercial for ExTran 7, and gives three case study synopses to support the product. Since the previous part of the paper is high quality, the commercial message can be tolerated, and the ad does provide some commercial state-of-the-art information. But listing machines and operating systems for which it is supported, contact addresses, etc., goes too far.

The use of inductive knowledge acquisition for different domains is the topic of the following three papers. Quinlan *et al* describe the derivation of rules for diagnosis of thyroid disease from archives of patient records. Then there is a report of an application in which rules for trouble-shooting a sinter plant (first stage in a lead smelting process) have been obtained from analyzing examples. This is thought to be the first expert system that was working at an industrial site in Australia. The third paper gives details of rule induction for classifying radar emitters, and presents some early experimental results. The final paper looks at problems posed by the construction of large knowledge bases. There is a discussion on distinctions between knowledge and data, and an argument for normalization of knowledge in similar fashion to normalization of data in databases. The argument is based on a rather narrow view of knowledge. I found the comments on the relationship of expert systems to conventional information systems more illuminating.

The average length of the papers is 18 pages, although the two by Winston and Waterman are longer. I particularly enjoyed the papers on legal decision making, window frame design, and thyroid data rule induction. They give a good view of the application areas and how particular techniques have been applied to them. Like most of the papers in the book, they describe the systems clearly, and discuss limitations and problems too.

Overall, the collection of papers is tightly prepared and well presented. There is a uniform style throughout the book. Careful editing has achieved this, although I did find one unmatched reference in the text. The book has an index. I would recommend the book as extremely useful for practitioners and students who want to understand some of the issues of expert systems in the outside world.

Dick Peacocke is manager of Knowledge Technology at BNR, President of CSCSI/SCEIO, and Adjunct Professor at Carleton University. He has a Ph.D. in Computer Science from University of Toronto.

*Hunt, E.B., Marin, J., and Stone, P. *Experiments in Induction*, New York, Academic Press, 1966.

Briefly Noted

What Every Engineer Should Know about Artificial Intelligence

Taylor, William A.

Cambridge, MA: The MIT Press, 1988, xi+331 pp
Hardbound, ISBN 0-262-20069-4, US$25.00

What it is that every engineer should know about AI is, apparently, a lot about the programming languages and software tools, quite a bit about AI as a social and industrial activity, and also something about inference systems and expert system applications. Little is said about computer vision or language. The style is breezy, and many of the frequent attempts at humour are successful. The title is unnecessarily restrained; the content is applicable not just to engineers, but to just about anyone who wonders whether this new-fangled AI stuff could be relevant to their own industrial problems. The main problem is that Taylor sees AI not as a science but as (in the words of the jacket blurb) "a few software ideas that work well enough to be used". For all its failures, there's more to AI than that.

—*G.H.*

University of Waterloo
Department of Computer Science

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

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

In accordance with Canadian Immigration requirements, this advertisement is directed to Canadian Citizens and Permanent Residents. The University encourages both women and men to apply.

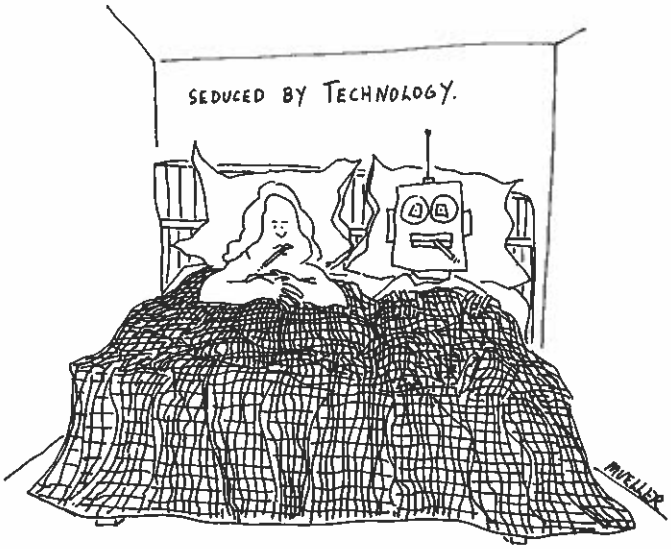

Careers in Artificial Intelligence, the magazine

The mission of the *Canadian Artificial Intelligence* magazine is to report on and promote AI research and development in Canada. We publish articles on active AI research in Canada and report on industrial AI development efforts. Included are book reviews, conference reports and a comprehensive list of technical reports and upcoming conferences.

Now is your chance to volunteer and join our dynamic team. Act now! Avoid the rush!

Positions:

- ***Surveyors/Interviewers***: Conduct a survey of AI graduate programs in Canada; talk to Profs and students—an excellent avenue for research in fault diagnosis, truth maintenance and knowledge acquisition.
- ***French translator***: Translate AI news into French.

Qualifications: If you think you can do it, you probably can.

Benefits: See your name in print. Fatten up the resume. Impress your friends.

All applicants will be entertained. Please send enquiries and replies, quoting No-Competition No. 1, to:

Canadian AI Magazine
Alberta Research Council
3rd Floor, 6815 - 8 Street N.E.
Calgary, Alberta
T2E 7H7
CDNnet: CSCSI@noah.arc.cdn
UUCP: cscsi%noah.arc.cdn@alberta.uucp

Books Received

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

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

Advances in Natural Language Generation: An Interdisciplinary Perspective
Zock, Michael; Sabah, Gérard (editors)
[LIMSI, Orsay]
(Communication in artificial intelligence series)
London: Pinter Publishers, 1988,
Two volumes: xix+200 pp and xviii+176 pp
Hardbound, ISBN 0-86187-965-1 and -995-5,
UK£27.50 per volume

Aspects of Speech Technology
Jack , Mervyn A.; Laver, John (editors)
[Centre for Speech Technology Research, University of Edinburgh]
(Edinburgh Information Technology Series 4)
Edinburgh University Press, 1988, viii+291 pp
[Distributed in the U.S. by Columbia University Press]
Hardbound, ISBN 0-85224-568-8, US$55.00

* **On Being a Machine**
Volume I: Formal Aspects of Artificial Intelligence
Narayanan, A.
[University of Exeter]
(Ellis Horwood series in artificial intelligence foundations and concepts)
Chichester: Ellis Horwood, 1988, 200 pp
(Distributed in Canada by John Wiley & Sons Canada)
Hardbound, ISBN 0-85312-957-6, Cdn$55.50

‡ **Cognizers: Neural Networks and Machines that Think**
Johnson, R. Colin; Brown, Chappell
(Wiley Science Editions)
John Wiley and Sons Canada Ltd, 1988, xi+260 pp
Hardbound, ISBN 0-471-61161-1, Cdn$32.95

Computers and Languages: Theory and Practice
Nijholt, Anton
[Free University, Brussels]
(Studies in computer science and artificial intelligence 4)
Amsterdam: North-Holland, 1988, xiii+482 pp
Hardbound, ISBN 0-444-70463-9, US$89.50

‡ **Exploring Artificial Intelligence: Survey Talks from the National Conferences on Artificial Intelligence**
Shrobe, Howard E. (editor)
[Symbolics Inc]
Morgan Kaufmann Publishers, 1988, xii+693 pp
(Distributed in Canada by John Wiley & Sons Canada)
Paperback, ISBN 0-934613-67-2, US$19.95;
Hardbound, ISBN 0-934613-69-9, US$39.95

* **Foundations of Deductive Databases and Logic Programming**
Minker, Jack (editor)
[University of Maryland]
San Mateo, CA: Morgan Kaufmann, 1988, 746 pp
Paperbound, ISBN 0-934613-40-0, US$36.95

Integrated Natural Language Dialogue: A Computational Model
Frederking, Robert E.
[Siemens, A.G.]
(The Kluwer international series in engineering and computer science; natural language processing and machine translation)
Boston: Kluwer Academic Publishers, 1988, xvii+178 pp
Hardbound, ISBN 0-89838-255-6, US$39.95

Lexical Ambiguity Resolution: Perspectives from Psycholinguistics, Neuropsychology, and Artificial Intelligence
Small, Steven L.; Cottrell, Garrison W.; Tanenhaus, Michael K. (editors)
[University of Rochester; University of California at San Diego; University of Rochester]
San Mateo, CA: Morgan Kaufmann, 1988, ix+518 pp
(Distributed in Canada by John Wiley & Sons Canada)
Hardbound, ISBN 0-934613-50-8, US$34.95

* **Practical Planning: Extending the Classical AI Planning Paradigm**
Wilkins, David E.
[SRI International]
(The Morgan Kaufmann series in representation and reasoning)
San Mateo: Morgan Kaufmann, 1988, xiii+205 pp
Hardbound, ISBN 0-934613-94-X, US$34.95

* **Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference**
Pearl, Judea
[UCLA]
(The Morgan Kaufmann series in representation and reasoning)
San Mateo: Morgan Kaufmann, 1988, xix+552 pp
Hardbound, ISBN 0-934613-73-7, US$39.95

Proceedings: Case-based Reasoning Workshop
Kolodner, Janet (editor)
[Georgia Institute of Technology]
Morgan Kaufmann Publishers, 1988, xiii+482 pp
(Distributed in Canada by John Wiley & Sons Canada)
Paperback, ISBN 0-934613-93-1, Cdn$42.95

Prosody and Speech Recognition
Waibel, Alex
[Carnegie Mellon University]
(Research notes in artificial intelligence)
San Mateo: Morgan Kaufmann
and London: Pitman, 1988, xii+212 pp
Paperback, ISBN 0-934613-70-2 and 0-273-08787-8, US$22.95

* **Reasoning with Incomplete Information**
Etherington, David

[AT&T Bell Laboratories]
(Research notes in artificial intelligence)
San Mateo, CA: Morgan Kaufmann
and London: Pitman, 1988, viii+240 pp
Paperbound, ISBN 0-934613-60-5, and 0-273-08785-1, US$22.95

‡ **Rule-based Programming with OPS5**
Cooper, Thomas A.; Wogrin, Nancy
[Digital Equipment Corporation]
Morgan Kaufmann Publishers, 1988, xiv+350 pp
(Distributed in Canada by John Wiley & Sons Canada)
Hardbound, ISBN 0-934613-51-6, Cdn$55.50

Semantic Networks: An Evidential Formalization and its Connectionist Realization
Shastri, Lokendra
[University of Pennsylvania]
(Research notes in artificial intelligence)
San Mateo, CA: Morgan Kaufmann
and London: Pitman, 1988, x+222 pp
Paperbound, ISBN 0-934613-39-7 and 0-273-08779-7, US$22.95

From Syntax to Semantics: Insights from Machine Translation
Steiner, Erich; Schmidt, Paul; Zelinsky-Wibbelt, Cornelia (editors)
[Institut für Angewandte Informationsforschung, Saarbrücken]
(Communication in artificial intelligence series)
London: Pinter Publishers, 1988, vii+262 pp
Hardbound, ISBN 0-86187-960-0, UK£29.50

‡ **Understanding Cognitive Science**
McTear, Michael F.
[University of Ulster at Jordanstown]
Chichester: Ellis Horwood, 1988, 264 pp
(Ellis Horwood series in cognitive science)
(Distributed in Canada by John Wiley & Sons Canada)
Hardbound, ISBN 0-7458-0161-7, Cdn$75.95

Computational Intelligence

Abstracts for 4(3) August 1988

Representation and Combination of Uncertainty with Belief Functions and Possible Measures
Didier Dubois and Henri Prade
Laboratoire Langages et Systemes Informatiques
Universite Paul Sabatier
Toulouse, France

The theory of evidence proposed by G. Schafer is gaining more and more acceptance in the field of artificial intelligence, for the purpose of managing uncertainty in knowledge bases. One of the crucial problems is the one of combining uncertain pieces of evidence stemming from several sources, whether rules or physical sensors. This paper examines the framework of belief functions in terms of expressive power for knowledge representation. It is recalled that probability theory and Zadeh's theory of possibility are mathematically encompassed by the theory of evidence, as far as the evaluation of belief is concerned. Empirical and axiomatic foundations of belief functions and possibility measures are investigated. Then the general problem of combining uncertain evidence is addressed, with focus on the Dempster rule of combination. It is pointed out that this rule is not very well adapted to the pooling of conflicting information. Alternative rules are proposed to cope with this problem, and deal with specific cases such as non-reliable sources, non-exhaustive sources, inconsistent sources, and dependent sources. It is also indicated that combination rules issued from fuzzy set and possibility theory look more flexible than the Dempster rule, because many variants exist, and their numerical stability seems to be better.

Multivalued Logics: A Uniform Approach to Reasoning in Artificial Intelligence
Matthew L. Ginsberg
Department of Computer Science
Stanford University
Stanford, California

This paper describes a uniform formalization of much of the current work in AI on inference systems. We show that many of these systems, including first-order theorem provers, assumption-based truth maintenance systems (ATMS's) and unimplemented formal systems such as default logic or circumscription can be subsumed under a single general framework.

We begin by defining this framework, which is based on a mathematical structure known as "bilattice". We present a formal definition of inference using this structure, and show that this definition generalizes work involving ATMS's and some simple nonmonotonic logics.

Following the theoretical description, we describe a constructive approach to inference in this setting; the resulting generalization of both conventional inference and ATMS's is achieved without incurring any substantial computational overhead. We show that our approach can also be used to implement a default reasoner, and discuss a combination of default and ATMS methods that enables us to formally describe an "incremental" default reasoning system. This incremental system does not need to perform consistency checks before drawing tentative conclusions, but can instead adjust its beliefs when a default premise or conclusion is overturned in the face of convincing contradictory evidence. The system is therefore much more computationally viable than earlier approaches.

Finally, we discuss the implementation of our ideas. We begin by considering general issues that need to be addressed when implementing a multivalued approach such as that we are proposing, and then turn to specific examples showing the results of an existing implementation. This single implementation is used to solve a digital simulation task using first-order logic, a diagnostic task using ATMS's as suggested by de Kleer and Williams, a problem in default reasoning as in Reiter's default logic or McCarthy's circumscription, and to solve the same problem more efficiently by combining default methods with justification information. All of these applications use the same general-purpose bilattice theorem prover and differ only in the choice of bilattice being considered.

Reasoning About Knowledge and Belief: A Survey

Gregory L. McArthur
Department of Computer Science
University of Toronto
Toronto, Ontario

We examine a number of logics of knowledge and belief from the perspective of knowledge-based systems. We are concerned with the beliefs of a knowledge-based system, including both the system's base set of beliefs — those garnered directly from the world — and beliefs which follow from the base set. Three things to consider with such logics are the expressive power of the language of the logic, the correctness and completeness of the inferences sanctioned, and the speed with which it is possible to determine whether a given sentence is believed. The influential possible worlds approach to representing belief has the property of logical omniscience, which makes for inferences which are unacceptable in the context of belief and may take too much time to make. We examine a number of weak logics which attempt to deal with these problems. These logics divide into three categories — those which admit incomplete or inconsistent situations into their semantics, those which posit a number of distinct states for a believer which correspond roughly to frames of mind, and those which incorporate axioms or other syntactic entities directly into the semantics. As to expressive power, we consider whether belief should be represented by a predicate or a sentential operator, and examine the boundary between self-referential and inconsistent systems. Finally, we consider logics of "believing only", which add the assumption that a system's base set of beliefs are, in a certain sense, all that it believes.

Abstracts for 4(4) November 1988

Hierarchical Planning Involving Deadlines, Travel Time and Resources

Thomas Dean
Department of Computer Science
Brown University
R. James Firby
Department of Computer Science
Yale University
David Miller
Department of Computer Science
Virginia Polytechnic Institute & State University

This paper describes a planning architecture that supports a form of hierarchical planning well suited to applications involving deadlines, travel time, and resource considerations. The architecture is based upon a temporal database, an heuristic evaluator, and a decision procedure for refining partial plans. A partial plan consists of a set of tasks and constraints on their order, duration, and potential resource requirements. The temporal database records the partial plan on which the planner is currently working, and computes certain consequences of that information to be used in proposing methods to further refine the plan. The heuristic evaluator examines the space of linearized extensions of a given partial plan in order to reject plans that fail to satisfy basic requirements (e.g., hard deadlines and resource limitations), and estimates the utility of plans that meet those requirements. The information provided by the temporal database and the heuristic evaluator is combined using a decision procedure that determines how best to refine the current partial plan. Neither the temporal database nor the heuristic evaluator is complete, and, without reasonably accurate information concerning the possible resource requirements of the tasks in a partial plan, there is a significant risk of missing solutions. A specification language that serves to encode expectations concerning the duration and resource requirements of tasks greatly reduces this risk, enabling useful evaluations of partial plans. Details of the specification language and examples illustrating how such expectations are exploited in decision making are provided.

Exploiting Temporal Coherence in Nonlinear Plan Construction

Mark Drummond and Ken Currie
AI Applications Institute
University of Edinburgh
Edinburgh,UK

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

Localized Event-Based Reasoning for Multiagent Domains

Amy L. Lansky
Artificial Intelligence Center
SRI International
Menlo Park, California
and Center for Study of Language and Information
Stanford University
Stanford, California

This paper presents the GEM concurrency model and GEMPLAN, a multiagent planner based on this model. Unlike standard state-based AI representations, GEM is unique in its explicit emphasis on events and domain structure. In particular, a world domain is modeled as a set of regions composed of interrelated events. Event-based temporal logic constraints are then associated with each region to delimit legal domain

behaviour. The GEMPLAN planner directly reflects this emphasis on domain structure and constraints. It can be viewed as a general purpose constraint satisfaction facility which constructs a network of interrelated events (a "plan") that is subdivided into regions ("subplans"), satisfies all applicable regional constraints, and also achieves some stated goal.

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

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

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

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

This paper presents a method of solving planning problems that involve actions whose effects change according to the situations in which they are performed. The approach is an extension of the conventional planning methodology in which plans are constructed through an iterative process of scanning for goals that are not yet satisfied, inserting actions that produce the same effects in every situation. The extension involves introducing additional subgoals to actions above and beyond the preconditions of execution normally introduced. These additional subgoals, called secondary preconditions, ensure that the actions are performed in contexts conducive to producing the effects we desire. This paper defines and analyzes secondary preconditions from a mathematically rigorous standpoint and demonstrates how they can be derived from regression operators.

Plans and Resource-Bounded Practical Reasoning
Michael E. Bratman
Department of Philosophy
Stanford University
and Center for the Study of Language and Information
David J. Israel and Martha E. Pollack
Artificial Intelligence Center
SRI International
and Center for the Study of Language and Information

An architecture for a rational agent must allow for means-end reasoning, for the weighing of competing alternatives, and for interactions between these two forms of reasoning. Such an architecture must also address the problem of resource boundedness. We sketch a solution of the first problem that points the way to a solution of the second. In particular, we present a high-level specification of the practical-reasoning component of an architecture for a resource-bounded rational agent. In this architecture, a major role of the agent's plans is to constrain the amount of further practical reasoning she must perform.

Included in V4 (4) will be two follow up papers to the "Taking Issue" section on "Forum: A Critique of Pure Reason", edited by Hector Levesque, which appeared in V3 (3). These papers are entitled:

FIAT LUX: A Reply to McDermott's "A Critique of Pure Reason"
Tim Flannagan
Logica Cambridge Ltd.
Cambridge, England

NO ABSTRACT

Exhuming the Criticism of the Logicist
Randy Goebel
Department of Computing Science
University of Alberta
Edmonton, Alberta

McDermott has recently explained his fundamental philosophical shift on the methodology of artificial intelligence, and has further suggested that the shift is both necessary and inevitable. The shift results from a perception that a trend towards over-formalization has detached the real problems from the research results. McDermott's criticism is an enlightened exhumation of the criticisms of the seventies, and explains new ways in which the logical methodology can be abused. I argue that McDermott's criticism should not discourage the use of logic, but force a timely re-examination of its fundamental role in AI.

Technical Reports

University of Alberta

Representational Difficulties with Classifier Systems
Dale Schuurmans and Jonathan Schaeffer
Technical Report TR-88-10

Classifier systems are currently in vogue as a way of using genetic algorithms to demonstrate machine learning. However, there are a number of difficulties with the formalization. These problems influence how knowledge is represented and the rate at which the system can learn. Some of the problems are inherent in classifier systems, and one must learn to cope with them, while others are pitfalls waiting to catch the unsuspecting implementer. Most of these problems can be illustrated with the game of tic-tac-toe. That this simple problem shows up many of the weaknesses of classifier systems suggests that a more powerful representation scheme is required.

Mizar-MSE Primer and User's Guide
edited by *P. Rudnicki*
Technical Report TR-88-11

The language Mizar-MSE is a formal language for recording mathematical texts of multisorted first order predicate calculus with equality. The language has proven convenient both notationally and for computer processing. The latter consists of checking logical validity of Mizar-MSE texts by computer processors for this language; they are implemented on several types of computers.

Handling Scope Ambiguities Using Domain-Independent Heuristics
Sven O. Hurum
Technical Report TR-88-12

The aim of this work has been to extend a general purpose natural language understanding system currently being developed in this department to handle "scope ambiguities". Scope ambiguities result from the interaction of logical operators such as quantifiers, coordinators and negation and are designated as such because the ambiguities are usually represented in terms of the relative scopes of the operators.

Particular attention is given in this study to the problems associated with the logical representation and scoping of indefinite noun phrases. Evidence is presented that giving indefinites a "referential" interpretation, as has been recently proposed, cannot account for their immunity to certain constraints on scoping. Rather, it is argued that a major distinction needs to be made between the scoping of existential and distributive quantifiers.

A scoping algorithm is presented which generates the set of valid scoped readings of a sentence, omitting logically redundant readings, and places these in an approximate order of preference using a set of domain-independent heuristics. The heuristics are based on information about the lexical type of each operator and on structural relations between pairs of operators. A discussion is given of the results of running the program on some test data and the report closes with some suggestions for possible extensions to the current program.

Nonmonotonic Reasoning in Temporal Domains: The Knowledge Independence Problem
Scott D. Goodwin and Randy G. Goebel
Technical Report TR-88-17

Much interest has been focused on nonmonotonic reasoning in temporal domains since Hanks and McDermott discovered that intuitive temporal representations give rise to the multiple extension problem. Here we consider nonmonotonic reasoning in temporal domains from the perspective of the Theorist hypothetical reasoning framework. We show how this framework can be applied to temporal reasoning in a simple and intuitive way to solve many of the problems posed in the recent literature, such as the Yale Shooting problem, Kautz's Vanishing Car problem, Haugh's Assassin problem and Haugh's Robot problem.

The basis of our solution to these problems is the characterization of the notion that the past is independent of the future (*temporal independence*) and the provision of two additional modes of reasoning: *conditional* explanation and prediction. The problem of representing and reasoning about temporal independence is an instance of a more general problem which we call the *knowledge independence problem*. In this paper, we provide a preliminary definition of the knowledge independence problem; we leave to future work further development of the obvious connections with statistical independence. Using our preliminary definition, we show how to represent and reason about temporal independence and how this solves many temporal reasoning problems.

— *Other relevant reports*

Statistically Founded Degrees of Belief
Fahiem Bacchus
Technical Report TR-87-2a

Experiments in Distributed Game-Tree Searching
Jonathan Schaeffer
Technical Report TR-87-2b

Three Papers on the Logical Form of Mass Terms, Generics, Bare Plurals, and Habituals
Francis Jeffry Pelletier and Lenhart K. Schubert
Technical Report TR-87-3

Predicates with Probabilities
Fahiem Bacchus and Lenhart Schubert
Technical Report TR-87-5

Searching for Chess
T.A. Marsland
Technical Report TR-87-6

Overheads in Loosely Coupled Parallel Search
E. Altmann, T. Breitkreutz, and T.A. Marsland
Technical Report TR-87-15

Ordering information

Requests for free copies of any of the forenamed publications should be addressed to: Computing Science Reading Room, 604 General Services Building, Department of Computing Science, University of Alberta, Edmonton, Alberta, Canada T6G 2H1;
Email: pinch@alberta.uucp.

University of Calgary

Directing the User Interface: How People Use Command-based Computer Systems
S. Greenberg and I.H. Witten
88/289/01, January 1988
Published in *Proc 3rd IFAC Conference on Man-Machine Systems*, pp 299-305, June 1988

Several striking and surprising characteristics of how people use interactive systems are abstracted from a large body of recorded usage data. In particular, we examine frequencies of invocation of commands and complete command lines (including modifiers and arguments), as well as vocabulary growth. Individual differences are of particular interest, and the results are analyzed by the user and by identifying groups of like users. The study underlines the remarkable diversity that exists even within groups having apparently similar needs.

Connecting to the Past
B.A. MacDonald
88/290/02, January 1988
Published in the collected papers of *IEEE Conference on Neural Information Processing Systems*, November 1987

Recently there has been renewed interest in neural-like processing systems, evidenced for example in the two volumes *Parallel Distributed Processing* edited by Rumelhart and McClelland, and discussed as parallel distributed systems, connectionist models, neural nets, value passing systems and multiple context systems. Dissatisfaction with symbolic manipulation paradigms for artificial intelligence seems partly responsible for this attention, encouraged by the promise of massively parallel systems implemented in hardware. This paper relates simple neural-like systems based on multiple context to some other well-known formalisms — namely production systems, k-length sequence prediction, finite-state machines and Turing machines — and presents earlier sequence prediction results in a new light.

On Asking the Right Questions
B.J. Krawchuk and I.H. Witten
88/293/05, January 1988
Published in abbreviated form in *Proc Fifth International Conference on Machine Learning*, pp 15-21, June 1988

Concept learning systems have a great deal to gain by showing more initiative; in particular, by actively exploring rather than passively waiting for more examples to appear. This paper briefly reviews different approaches to selecting questions, and goes on to explore the ramifications of one in detail. Current implementations are overly simplistic because they assume a hierarchically-structured network of concepts. It is shown how they break down in more general partially-ordered domains, and a new method which copes with this situation is described and illustrated. Finally, the techniques are related to the version-space approach to provide them with a well-understood theoretical underpinning.

Teaching a Mouse How to Draw
D.L. Maulsby and I.H. Witten
88/294/06, January 1988

A scheme is described for graphical programming by example, which lets untrained end-users add composite operations to a drawing program using constructive methods traditionally employed in drafting. A pilot experiment showed that considerable extraneous activity occurs in naturally-produced traces (errors, doodling, false starts, etc). To combat this, fullest advantage is taken of the interactive situation to constrain induction by suppressing, or at least controlling, gratuitous variation. A flatland device called "Meta-Mouse" serves to concentrate the user's attention on the job of teaching a student with limited capabilities. It predicts actions, asks for constructions, and solicits input parameters when required. It carries out anticipated actions promiscuously, to encourage consistency and force the teacher to satisfy appropriate felicity conditions. Implications for machine learning include the benefits of simulating a pupil to complete the teaching metaphor, and the positive role that close user interaction can play in constraining the search for apt generalizations.

Inducing Functions in Robot Domains
D. Pauli and B.A. MacDonald
88/296/08, January 1988

This paper presents a reimplementation of a function induction algorithm, representing functions in logic rather than the original nested expressions. The algorithm is part of Peter Andreae's robot procedure learning system Noddy. Results are given, the implementation is compared to Andreae's version, and also to related work; the Bacon, Coper, Abacus, Marvin, Basil and Pengi systems. The paper focuses on the representation of knowledge for inducing functional relations — in particular on the representation of inverse operators and the induced expression — and the strength of argument typing.

Acquiring Graphical Know-How: an Apprenticeship Model
D.L. Maulsby and I.H. Witten
88/302/14, March 1988
Published in *Proc European Knowledge Acquisition Workshop 88*, pp 34/1-34/16, June 1988

This paper studies the acquisition of procedural knowledge, or "know-how", from end-users in the domain of interactive graphics. In order to develop an open-ended system that is not restricted to any particular class of drawings, heavy emphasis is placed on the user interface. Experts (we call them simply "teachers") express procedures constructively, using any of the tools available in the interactive drawing environment. Well-structured procedures, including branching and looping, are inferred using a variety of weak generalization heuristics. The teacher's attention is concentrated on the system's perceptual and inferential shortcomings through a metaphorical apprentice called "Meta-Mouse". Its sensors are predominantly tactile, which forces

teachers to make their constructions explicit. Meta-Mouse generalizes action sequences on the fly and eagerly carries out any actions it can predict. Theoretical support for the design comes from two sources: geometric phenomenology, which confirms that powerful problem-solving methods are associated with spatial reasoning; and the fact that Meta-Mouse automatically imposes important "felicity conditions" on the teacher's demonstrations.

A Framework for Knowledge Acquisition through Techniques of Concept Learning
B.A. MacDonald and I.H. Witten
88/303/15, March 1988
To be published in *IEEE Trans Systems, Man and Cybernetics*, special issue on knowledge engineering

Knowledge-based systems must represent information abstractly so that it can be stored and manipulated effectively. Schemes for learning suitable representations — or concepts — from examples promise domain experts direct interaction with machines to transfer their knowledge. This paper develops an integrative framework for describing concept learning techniques which enables their relevance to knowledge engineering to be evaluated. The framework provides a general basis for relating concept learning to knowledge acquisition, and is a starting point for the development of formal design rules.

The paper first frames concept learning in the context of knowledge acquisition. It then discusses the general forms of input and concept representation: as *logic*, *functions* and *procedures*. Next, methods of biasing the search for a suitable concept are described and illustrated: *background knowledge*, *conceptual bias*, *composition bias* and *preference orderings*. Then modes of teacher interaction are reviewed: the *nature of examples given*, and the *method of presenting them*. Finally, the framework is illustrated by applying it to the better-documented concept learning systems.

I, Metamouse
I.H. Witten and D.L. Maulsby
88/306/18, June 1988
Published in *Proc Symposium "10 Years IIG"*, pp 198-210, June 1988

"Metamouse" is a Flatland device that learns procedures in interactive graphics taught to it by untrained users. It lives in a simple graphics world and induces procedures in real-time, taking full advantage of the interactive situation to suppress, or at least control, variability in the teaching sequence. The teacher must demonstrate what is to be done, employing constructive methods to make constraints readily apparent. Metamouse predicts actions, asks for constructions, solicits input parameters when required and induces a program (including conditionals and loops). By eagerly carrying out predicted actions (which the teacher can undo if they are incorrect) it helps the teacher to satisfy appropriate felicity conditions. In this paper, Metamouse explains its world, its body and mind, and what it can do.

Explanation-based Learning: its Role in Problem-Solving
B.J. Krawchuk and I.H. Witten
88/307/19, June 1988
To be published in *Journal of Experimental and Theoretical Artificial Intelligence*

"Explanation-based" learning is a semantically-driven, knowledge-intensive paradigm for machine learning which contrasts sharply with syntactic or "similarity-based" approaches. This paper redevelops the foundations of EBL from the perspective of problem-solving. Viewed in this light, the technique is revealed as a simple modification to an inference engine which gives it the ability to generalize the conditions under which the solution to a particular problem holds. We show how to embed generalization invisibly within the problem solver, so that it is accomplished as inference proceeds rather than as a separate step. The approach is also extended to the more complex domain of planning, which involves maintaining and operating on a global world state, to illustrate that it is by no means restricted to toy problem-solvers. We argue against the current trend to isolate learning from other activities and study it separately, preferring instead to integrate it into the very heart of problem solving.

Report on Function Induction
D.G. Pauli
88/308/20, June 1988

This report outlines research work from May 1 to August 31, 1987. It describes a function induction algorithm that is part of Peter Andreae's robot procedure learning system Noddy. The report also describes a C-Prolog implementation of this algorithm.

Two other systems that induce functions are Bacon and Coper. Both are discovery systems that determine empirical laws from observational data. In addition to these systems, a general approach to symbolic function induction is discussed.

This work was motivated by a desire to investigate the ramifications of using a non-functional language such as Prolog to do function induction. The discussion focuses on the representation of knowledge — in particular on the representation of inverse operators and the induced expression — and the strength of argument typing.

An Interactive Programming Environment for the Development of Expert Systems in Well Log Analysis
B.A. Lowden
88/309/21, June 1988

The job of a well log analyst is to infer properties of subsurface rock layers from measurements taken in a petroleum well bore. Although there are many Fortran programs for the numerical analysis and display of this data, log analysis also requires symbolic processing, and can benefit from the use of expert systems. This paper describes Interlog, an interactive programming environment designed to facilitate the development and use of expert systems in this field.

Interlog provides interpreters for Lisp, a Fortran-like language and an object-oriented language; a screen editor and other programming tools; Fortran libraries for graphics and data storage; and libraries for the management and analysis of well log data. These tools can be used to build expert systems which contain modules written in various styles (e.g., procedural or object-oriented), and integrate symbolic and numeric processing. On the other hand, in order to be most useful to log analysts, these expert systems must run on currently used hardware configurations and utilize existing software and data. Therefore, while Interlog supports Lisp, it is portable, can run under time-sharing systems and can incorporate Fortran software. Because expert systems developed with Interlog can fit so tightly with current log analysis software, they should have immediate commercial uses.

Modeling and Generating Music Using Multiple Viewpoints
D. Conklin and J.G. Cleary
88/313/25, July 1988
Published in *Proc First AAAI workshop on Artificial Intelligence and Music*, pp 125-137, August 1988

A technique for modeling tonal music using multiple viewpoints is described. Markov models are used with a different model for each of a number of different viewpoints: for example, the durations of notes or their relative intervals. The models are extracted from existing pieces, or can be constructed by hand. The technique is evaluated by using models to generate new pieces. Examples are given of pieces generated using models of monodic Gregorian chant and two voice polyphony.

The Other Side of the Coin: Teaching Artificial Learning Systems
I.H. Witten and B.A. MacDonald
88/318/30, August 1988

The burgeoning technology of machine learning is beginning to provide some insight into the nature of learning and the role of teaching in expediting the learning process. A number of systems that learn concepts and procedures from examples have been described in the research literature. In general these require a teacher who not only has an analytical understanding of the problem domain, but also is familiar with some of the internal workings of the learning system itself. This is because the learner is performing a search in concept space which is generally quite intractable, but for the teacher's selection of guiding examples.

A concept learning system's teacher must select a complete, properly ordered set of examples — one that results in a successful search by the system for an appropriate concept description. In some systems, the set of examples determines whether the concept can or cannot be learned, while the order of presentation affects execution time alone. In others, both examples and presentation order are jointly responsible for success. Yet others occasionally select critical examples themselves and present them to a teacher for classification. In all cases, however, the teacher provides the primary means whereby search is pruned. Sometimes the teacher must prime the learner with considerable initial knowledge before learning can begin.

Not surprisingly, systems which demand more of the teacher are able to learn more sophisticated concepts. This paper examines the relationship between teaching requirements and learning power for current concept learning systems. We introduce concept learning by machine with emphasis on the role of the human teacher in rendering practical an otherwise intractable concept search. Machine learning has drawn many lessons from human learning and will continue to do so. In turn it can contribute more formal, if simpler, analyses of concept learning from examples.

Of Mice and Pens: Human Performance in Drawing
U.Y. Chow, D.L. Maulsby and I.H. Witten
88/319/31, September 1988

When asked to draw with pen on paper, people exhibit surprising regularity in the apparently free choices they make to execute primitive strokes. Some patterns can be explained in terms of the mechanics of holding the writing instrument; others stem from economy of motion; yet others signify preferred ways of achieving precision when anchoring lines. This paper describes a series of experiments designed to test the extent to which the effects carry over to drawing with mouse and drafting program. It concludes that some habits transfer, albeit in weaker form, despite the fact that mechanical constraints are radically different.

Inducing Programs in a Direct-Manipulation Environment
D.L. Maulsby and I.H. Witten
88/320/32, September 1988

End users who need to program within highly interactive direct-manipulation interfaces should be able to communicate their intentions through concrete demonstration rather than in terms of symbolic abstractions. This paper describes a system that learns procedures in interactive graphics taught to it "by example" by minimally trained users. It shows how techniques of machine learning and reactive interfaces can support one another — the former providing generalization heuristics to identify constraints implicit in user actions, the latter offering immediate feedback to help the user clarify hidden constraints and correct errors before they are planted into the procedure. The teacher's attention is focused on the learning system's perceptual and inferential shortcomings through a metaphorical apprentice called Metamouse, which generalizes action sequences on the fly and eagerly carries out any actions it can predict. The success of the induction process is assessed quantitatively by counting how many erroneous predictions are made during example tasks.

Kinematics of an Elbow Manipulator with Forearm Rotation: the Excalibur
R. Heise and B.A. MacDonald
88/324/36, October 1988

General methods and typical specific solutions for robot arm geometry are well-known. This paper presents a detailed solution for a six joint manipulator which has a rotation at mid-forearm rather than a third wrist axis. Details are given indicating how real joint angles relate to those modeled by the more abstract kinematics. All degeneracies are considered and methods for handling

them are given. The paper provides a complete tutorial for kinematic modeling with a specific arm.

Autonomy, Intelligence, and Instructability
B.A. MacDonald and I.H. Witten
88/335/37, October 1988

Instructable systems constitute an important, useful and practically realizable step towards fully autonomous ones. In many applications people will not want machines to be self-motivated, but they will want to teach them new jobs. The user interface must permit the teacher to guide the system through tasks. The system employs samples of behavior so gathered to drive an inductive process of concept learning. Learning becomes intractable unless the teacher fulfills certain felicity conditions. The real world frequently constitutes a competitive learning environment, and instructable systems may have to guard against their knowledge and skills being corrupted by incorrect or deliberately misleading teachers. Experimental prototypes of two instructable systems are presented, one for verbally editing robot movements, the other for automating office tasks. These examples show the potential utility of approaching autonomy via instructability; the next steps are to extend the power of their learning mechanisms, and to render them robust.

Source Models for Natural Language
I.H. Witten and T.C. Bell
88/326/38, November 1988

A model of natural language is a collection of information that approximates the statistics and structure of the language being modeled. The purpose of the model may be to give insight into rules which govern how text is generated, or to predict properties of future samples of the language. This paper studies models of natural language from three different, but related, viewpoints. First, we examine the statistical regularities that are found empirically, based on the natural units of words and letters. Second, we study theoretical models of language, including simple random generative models of letters and words whose output, like genuine natural language, obeys Zipf's law. Innovation in text is also considered by modeling the appearance of previously unseen words as a Poisson process. Finally, we review experiments that estimate the information content inherent in natural text.

Simon Fraser University

— *School of Computing Science*

Specularity Removal for Shape-from-Shading
F. Tong
Technical Report CMPT TR 88-1
(M.Sc. thesis; supervisor: Funt)

Specularity reflecting surfaces confuse traditional shape-from-shading algorithms because the variation in image intensity within a specularity does not directly relate to the cosine of the incident angle, as it would for a simple Lambertian reflector. To overcome this problem, color is introduced and a method of removing the specular component of the intensity variation is proposed based on a dichromatic model of surface reflection. Unlike Shafer's method for specularity removal, which is restricted to uniformly colored surface patches, our algorithm uses information from several differently colored regions. The problem of segmenting an image into color regions is successfully avoided as the specular component is calculated and removed using local operations only. The image resulting from specularity removal preserves the relative intensity of the diffuse component so it can then be input to a shape-from-shading algorithm. Our shape-from-shading algorithm is based on variational calculus. Without assuming the location of the scene illuminant, and allowing background illumination, the algorithm computes the shape from the diffuse component image in a more general setting than the existing algorithms do. In the thesis, the algorithm is formulated into a local relaxation scheme which allows a parallel network implementation.

Ordering information
Requests for free copies of the forenamed publications should be addressed to: Technical Reports Requests, School of Computing Science, Simon Fraser University, Burnaby, British Columbia V5A 1S6.

— *Laboratory for Computer and Communication Research (LCCR)*

Single- and Multi-Chain Recursion: the Core of General Linear Recursion
Jiawei Han
Technical Report LCCR TR 88-3

We demonstrate that all the function-free recursive clusters consisting of a single linear recursive rule and one or more exit rules can be compiled into one of the three primitive classes: (1) bounded recursion, (2) single-chain recursion, or (3) multi-chain recursion, where a single- or multi-chain recursion is the recursion whose compiled formula contains a single compiled chain or multiple synchronized compiled chains respectively. Our study is based on a classification of linear recursion and a study of the compilation results of each class. Using a variable connection graph, we classify linear recursive clusters into six classes: acyclic paths, unit cycles, uniform cycles, nonuniform cycles, connected components and their disjoint mixtures and show that each class is compilable to bounded, single- or multi-chain recursion. The simplicity of our results is that a linear recursive cluster, no matter how complex its variable pattern is, is just a simple compiled formula in the processing point of view. The importance of our results is that it points out that the key to solving complex linear recursions is the efficient processing of single- and multi-chain recursions. Our study also develops a method of compiling and optimizing complex linear recursions using variable connection graphs.

A Semantic Basis for Explicit Belief: Preliminary Report
James P. Delgrande
Technical Report LCCR TR 88-5

A general framework for the investigation of logical systems of belief that are both tractable and semantically

well-motivated is presented. The approach extends standard possible worlds semantics in two ways. First, partial possible worlds, or situations, are employed. Second, the set of situations used to determine the truth of an explicit belief, $B\,a$, at a situation depends in part on the proposition expressed by a. It is argued that the approach provides an intuitive semantic basis from which systems may be constructed, contrasted, and compared. Proofs of soundness and completeness are given directly in terms of this semantics and not, as is the case with previous work, by appealing to similar results in relevance logic. Thus the formal results also provide a connection between the semantic theory of possible worlds and that of relevance logic.

Given this framework, we propose a specific system, BRIPK, as a "preferred" model of explicit belief. This system arguably retains a strong intuitive basis, while avoiding the (perceived) pitfalls of earlier systems. Moreover it is tractable and permits iterated modalities.

Truth Conditions and Procedural Semantics
Robert F. Hadley
Technical Report LCCR TR 88-6

On a common usage of "information", the information conventionally expressed by a declarative sentence is identical to its truth conditions. A sentence's truth conditions, in turn, have often been taken to be determined by, or to be identical with, its meaning. Meanings, in turn, are commonly explicated in terms of some model-theoretic semantics, e.g., the intensional semantics of Montague or Lewis, or the situation based semantics of Barwise and Perry. Such model-theoretic systems invariably postulate the existence of sets or collections whose cardinalities are, if not infinite, at least very large. The alleged existence of such sets (collections) raises difficult problems of set individuation which are, by now, familiar to many model theorists. It is argued herein that the solution to these individuation problems requires a deeper form of semantic theory, namely some form of Procedural Semantics (PS). Loosely construed, PS is the thesis that the meaning of a symbolic expression is identical to (or may be represented by) procedures which specify how the expression is to be used, or applied to the world. J.A. Fodor has characterized PS as a wildly implausible form of the verification theory of meaning. He has argued that PS constitutes a plausible semantic theory only for highly simplistic universes, such as "blocks worlds" and databases.

A version of PS is presented herein which avoids the simplistic "verification procedure" approach of early PS. The new theory embraces aspects of Quine's pragmatism, and assumes that semantic procedures may return pragmatic (default) truth values which may be revised, if the need arises, by adjudication procedures. On the theory here described, semantic procedures do not constitute the *complete* meanings of most symbolic expressions, but *constrain* their truth conditions. However, this default-oriented theory does require that *some* symbols (which may not occur in our public (shared) language) are tied to experience by computable, terminating procedures. It is argued that such procedures must exist if there is to be a coherent foundation for such traditional notions as *denotations* and truth conditions.

Methods of Formalizing Common Sense
Howard J. Hamilton
Technical Report LCCR TR 88-8

Techniques can be derived from common sense and implemented in artificial intelligence (AI) programs. In present implementations, common sense is not completely captured. A first step towards implementing common sense is to formalize it. In this paper, methods of formalizing common sense are explained in detail and classified.

First, an operational definition of common sense is proposed that emphasizes its inclusion of both commonsense knowledge and commonsense reasoning. Next, several approaches to formalizing common sense are described and classified according to an original classification scheme. One way of classifying the methods of formalizing common sense that have been developed in AI is by whether they emphasize knowledge or reasoning. In a *knowledge-oriented approach*, the focus is on facts that are encoded and not on which language they are represented and not on the manner in which the knowledge is to be used. In a *reasoning-oriented approach*, emphasis is given to the reasoning strategy to be used and not on the facts to be encoded. The distinction between knowledge-oriented and reasoning-oriented is not the same as the distinction between declarative and procedural representations because the classification is of the methods used by researchers to obtain formalisms, not of the formalisms themselves.

Another way of classifying the methods of formalizing common sense is top-down or bottom-up. A *top-down approach* begins with an existing, detailed reasoning scheme or set of knowledge and then seeks the underlying lower-level model. The process is iterated on the lower-level models. A *bottom-up method* tries to create a simple, fundamental model directly, perhaps from introspection or personal observations. This usage of the terms "top-down" and "bottom-up" differs from their use in describing control structures; again, the method of obtaining the formalism is being classified, not the control structure employed in the resulting formalism.

By combining these two ways of classifying, four distinct types of methods can be distinguished: bottom-up knowledge-oriented, top-down knowledge-oriented, bottom-up reasoning-oriented, and top-down reasoning-oriented. The largest portion of this paper describes each of these methods, gives examples of each and lists some of their assumptions, advantages and disadvantages. Where appropriate, comparisons are made among the methods.

The Multi-Way Counting Method
Jiawei Han
Technical Report LCCR TR 88-10

The multi-way counting method represents a group of counting algorithms which generalize the counting method to complex linear recursions and versatile recursive queries. Since the counting method has been recognized as one of the best performing algorithms in the processing of "simple" recursive queries on "typical" linear recursions, its generalizations to different query forms and complex recursions have great potential to

out-perform other algorithms. Since many complex single linear recursions can be compiled to n-chain recursions, we first study the multi-way counting on primitive n-chain recursions and then extend the result to general (single) linear recursions. A quad-state variable binding analysis method is used to analyze query instantiations and inquiries on multi-chain recursive predicates, which leads to counting in multiple direction combinations, hence the name, the multi-way counting method. A comparison of our method with some other previously studied processing methods is also discussed in the paper.

Generalized Single-Chain Recursions
Jiawei Han
SFU, Burnaby, British Columbia
Wenyu Lu
Ohio State University, Columbus, Ohio
Technical Report LCCR TR 88-11

A generalized single-chain recursion (GSC), whose compiled formula consists of single or multiple asynchronous compiled chains, is a generalization of transitive closure and single-chain recursions (SCs). Some previous studies show that many single linear recursive rule clusters (SLs) can be compiled to SCs. This paper shows that many non-SL recursive clusters, which include multiple linear recursive rules, nonlinear recursive rules, mutually recursive rules and multiple levels of recursions, can be compiled to GSCs. Our study on the compilation of non-SL clusters to GSCs, the simplification of compiled formulas and the query processing of GSCs shows that GSCs are frequently encountered, which can be compiled to relatively simple compiled formulas and processed using the methods similar to transitive closure processing strategies.

A Process-Oriented, Intensional Model of Knowledge and Belief
Robert F. Hadley
CSS/LCCR TR 88-16

A process-oriented, intensional model of knowledge and belief is presented which permits the representation of nested epistemic attitudes, involving different agents and different times. Inferences among these representations are sanctioned by the model (hereafter called NIM, for *nested intensional model*). NIM is axiomatized in first-order predicate logic, and addresses the fact that an agent's beliefs are affected by numerous *empirical* conditions. NIM also expresses and sanctions inferences to the effect that an agent's beliefs may be described by distinct, intensionally equivalent formulas. The model is designed to capture "human-like" epistemic concepts, and successfully avoids both logical omniscience and (what Levesque dubs) the unduly restrictive syntactic approach. Moreover, NIM provides agents with a conceptual map, interrelating the concepts of knowledge and belief and a number of cognate concepts, such as "infers", "retracts", and "questions". Because the model builds upon the concept of an *intension*, whose degree of granularity is left open, the range of formulas which are permitted to express the *same belief* is left open to the user of NIM. However, a number of semantic theories are described which yield plausible, "human-like" notions of synonymy and "same belief".

Static Discontinuity Grammars for Government-Binding Theory
Veronica Dahl
CSS/LCCR TR 88-22

We describe a logic grammar formalism tailored to accommodate language processing applications of Government-Binding theory. It consists of productions that rewrite some explicit symbols in a parsing state, using shared substitutions, while skipping over any other symbols in that parsing state; and which retain type-0 power together with the representational simplicity of context-freedom. Thus, they can describe movement to (relating of) arbitrarily distant positions, while inducing parse *trees* (not graphs) with respect to which hierarchically characterized linguistic constraints can be readily described and enforced. The formalism's main feature — static discontinuity — was introduced by the author in 1986, both for grammars and for programs. In the latter context it gives rise to Discontinuous Logic, in which strings of terms in a proof state can be skipped while expanding other explicit, discontinuous terms in that state, with shared substitutions.

Collative Semantics: A Study in the Discrimination of Meaning
Dan Fass
CSS/LCCR TR 88-24

A semantic theory for natural language processing (NLP) called Collative Semantics (CS) is proposed. The theory, "pretheory" (main assumptions underpinning the theory), and "metatheory" (means for evaluating the theory) are outlined. CS rests on an analysis of meaning in natural language from which two key semantic phenomena are identified: lexical ambiguity and what are termed "semantic relations." Seven semantic relations are distinguished: literal, metonymic, metaphorical, anomalous, redundant, inconsistent and novel relations. CS has been implemented in a computer program called meta5, which is described. Meta5 analyzes sentences, discriminates the seven kinds of semantic relation between pairs of word senses in those sentences and resolves any lexical ambiguity. CS is compared against other general theories in NLP, notably Wilks' Preference Semantics and the body of work based on Schank's Conceptual Dependency. Some implications of and possible extensions to CS are discussed.

Meta-programming for Discontinuous Grammars
Veronica Dahl and Pierre Massicotte
CSS/LCCR TR 88-25

In this article we discuss meta-programming for discontinuous grammars — i.e., logic grammars in which productions can skip over unidentified strings of constituents called skips, and reposition them without analyzing them. We present an interpreter for a specific discontinuous grammar family in which skips are not allowed to move (the *static discontinuity* family), and we give an economical methodology for constructing and processing their parse history for the purpose of enforcing constraints which can be expressed in terms of node domination relationships.

Discontinuous Grammars
Veronica Dahl
CSS/LCCR TR 88-26

This article ties together several years of research on discontinuous grammars — logic grammars in which non-explicit sequences of symbols can be alluded to in the rules, and sometimes repositioned by them. After an introduction, we define them formally, present their background and provide intuitive insight into their use. Next, we examine several motivating arguments, both from formal and from natural language processing viewpoints, and we discuss the *static discontinuity* family of these grammars, in which a) the non-explicit strings are now allowed to move, and b) linguistic constraints specifically designed to suit, in particular, Government and Binding theory, can be defined modularly and statically in terms of node domination in parse trees, and are enforced dynamically. Finally, we discuss implementation issues, related work, and extensions.

Ordering information
Requests for free copies of any of the forenamed publications should be addressed to: Administrative Assistant, LCCR, Simon Fraser University, Burnaby, British Columbia V5A 1S6. Phone: (604) 291-4704. Email: ethel@cs.sfu.ca

— *AI Theses*

Using Modal Structures to Represent Extensions to Epistemic Logics
S.J. Hamilton
(M.Sc. thesis; supervisor: Delgrande)
(presented at CSCSI-88)

Kripke structures have been proposed as a semantic basis for modal logics of necessity and possibility. They consist of a set of states, informally interpreted as "possible worlds", and a binary accessibility relation between states. The primitive notion of a possible world in this context seems highly intuitive, since necessity can be interpreted as "truth in all possible worlds", and possibility as "truth in some possible world". However, modal logics have also been used to model the epistemic notions of knowledge and belief, where an agent at a particular world is said to "know" or "believe" a proposition if that proposition is true in all possible worlds compatible with its beliefs. In this context, it is not as obvious how to interpret a "possible world".

Modal structures have recently been introduced as a formally equivalent alternative to Kripke structures for modeling particular states of knowledge and belief. Modal structures consist of an infinite number of recursively defined levels, where each level contains the possible worlds that model an agent's meta-beliefs of a certain depth. For example, beliefs about the world are modeled at level 1 of a modal structure, and beliefs about beliefs about the world are modeled at level 2. Each modal structure corresponds to a single world of a Kripke structure and contains all the worlds that are accessible from that world in its levels. Modal structures are defined for the classical propositional epistemic logics S4 and S5.

Recently, the traditional possible worlds approach has been extended to model "explicit", or limited, belief with partial worlds, called situations, in an appropriately modified Kripke structure. In this thesis, I demonstrate how modal structures can replace Kripke structures to interpret three recent logics of explicit and implicit belief. I also extend modal structures to model a first-order predicate logic which includes quantifiers, equality and standard names. For each logic, I demonstrate the equivalence of the extended modal structure and the Kripke structure that originally provided the semantics for the logic. I discuss the advantages and disadvantages of using modal structures to model logics of knowledge and belief.

University of Toronto

Inductive Inference and Stability
Peter Turney
Dept. of Philosophy
doctoral dissertation, June 6, 1988

This thesis is concerned with developing a model of inductive inference of patterns in sets of numbers. In philosophy, this is known as the curve fitting problem, when the sets of numbers are sets of pairs of real numbers. In computer science, this is a standard way of framing the problem of inductive inference, when the sets of numbers are ordered sets of integers. A solution to the curve fitting problem should contribute to the understanding of scientific inference, since much of scientific inference consists of induction of patterns in numerical data.

It is proposed that a pattern or theory about a set of numbers must be both accurate and simple. The problem with this proposal is that there is no good definition of simplicity. It is suggested here that simplicity is stability. Stability is the ability to resist damage due to random accidents. This heuristic definition leads to two precise definitions. The first precise definition is a measure of the stability of directed graphs. This is useful for measuring the stability of algorithms, since algorithms are naturally described by directed graphs (flowcharts are directed graphs). The second precise definition is a measure of the stability of linear regression equations. This provides a solution to the curve fitting problem for linear regression equations.

Application and Order Forms are on Page 60

ADVANCED COMPUTER TECHNOLOGY FOR RESOURCE INDUSTRIES

* FORESTRY * ELECTRIC POWER * MINING * OIL *

Specialists in

Expert Systems * Data Acquisition

Speech Recognition

Now Shipping

Paravant Ruggedized Computers

Autometrics Power Monitor System

VRP Expert System Developer's Kit

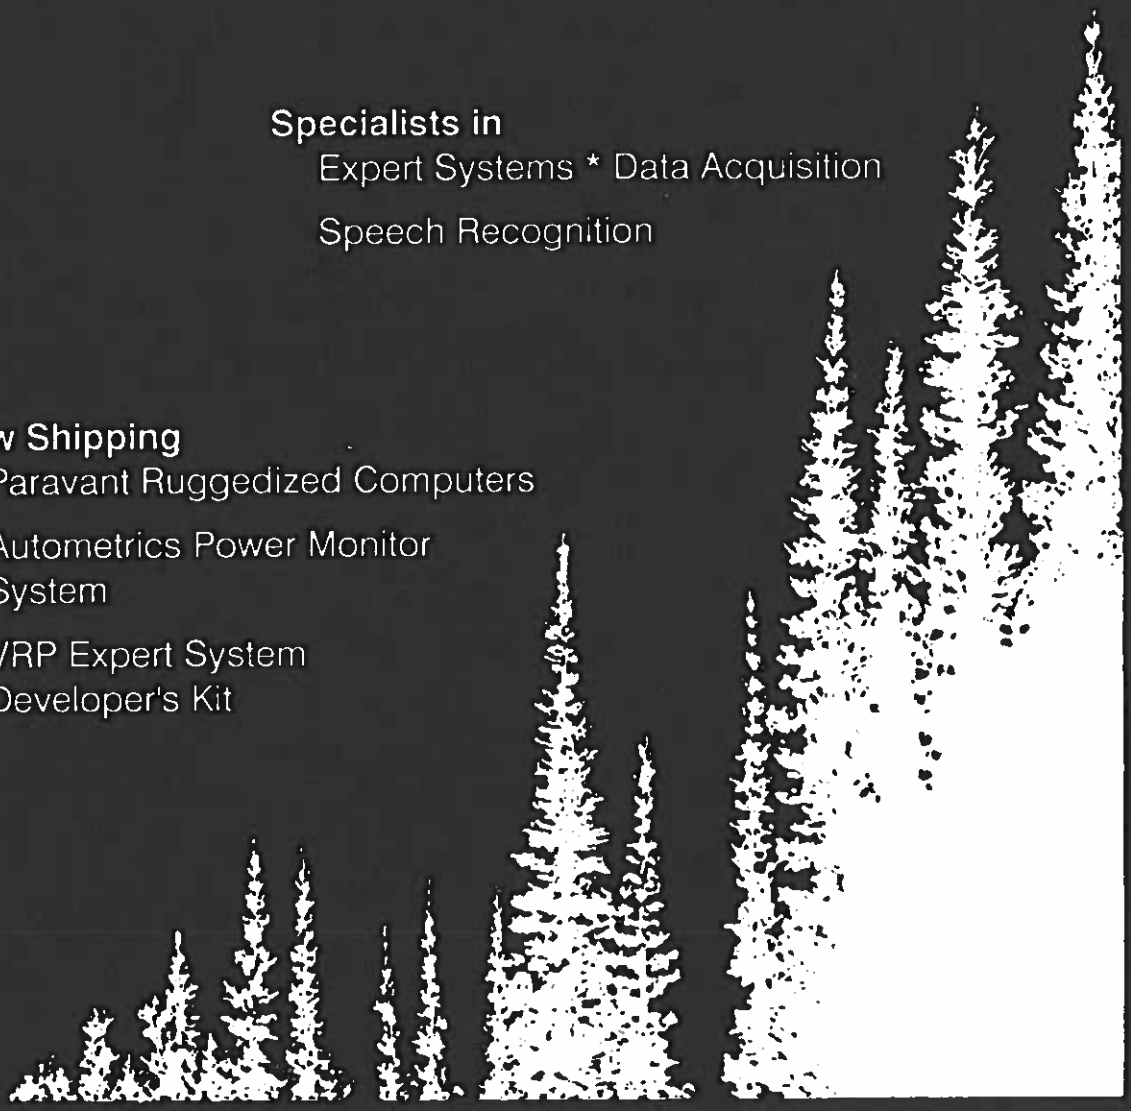

AUTOMETRICS

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

World Watch

World Watch on AI Applications and Development is sponsored by the National Research Council's (NRC) Associate committee on AI. Based on the information provided in the abstracts, the references provided have been selected by the secretariat of the NRC Associate Committee on AI as a representative sample of interest and value to Canadian industry. Abstracts provided are reprinted from "Key Abstracts in Artificial Intelligence" with permission from INSPEC. INSPEC is widely recognized as the leading English-language database covering the published information in the field of physics, electronics and computing. Information contained in the INSPEC services is collected on an international basis from over 4,000 Journals and 1,000 Conference Proceedings. INSPEC is a division of the Institution of Electrical Engineers, Station House, Nightingale Road, Hitchin, Herts, UK. All INSPEC's products and services are available in North America from the INSPEC Dept. IEEE Service Centre, 445 Hoes Lane, P. O. Box 1311, Piscataway, NJ 08855-1331, USA.

Persons wishing to obtain copies of references cited should contact their nearest technical library or the Canada Institute for Scientific and Technical Information (CISTI), NRC, Building M-55, Montreal Road, Ottawa, Ontario K2A 0S2, (Phone: (613) 993-1585, Telex: 053-03115). For on-line ordering, CAN/OLE users may use the CAN/DOC command. Envoy users type "COMPOSE CISTI."

1.0 THEORETICAL ASPECTS

1254 The Perceptron and AI: a New Old Way Forward?
D. Ellison
Dundee Coll. of Technology, UK
Comput. Bull. (UK), vol. 4, pt. 1, pp 27-30, March 1988

During the early days of artificial intelligence, Perceptrons were studied intensively. These "brain models" were loosely based on neural mechanisms and there was great hope that using the serial (Von Neumann architecture) digital computers of the day large leaps forward in AI would be made. Unfortunately, they went out of favour after several elementary limitations of Perceptrons were discovered. There is a growing resurgence in the "neural" based approach; as a result of improved experimental techniques, more is known about the detailed structure and functioning of neurons and brains; parallel computers are becoming generally available; and with the VLSI technology, it is feasible to put something akin to simple neural nets on chips at reasonable cost; and, finally, the mathematical theories of learning and distributed memory and processing have improved. (18 refs.)

1262 Expressiveness and Tractability in Knowledge Representation and Reasoning
H.J. Levesque
Dept. of Computer Science, U. of Toronto
Toronto, Ontario, Canada
R. J. Brachman
Computational Intelligence (Canada), vol.3, no.2
pp 78-93, May 1987

A fundamental computational limit on automated reasoning and its effect on knowledge representation is examined. Basically, the problem is that it can be more difficult to reason correctly with one representational language than with another and, moreover, that this difficulty increases dramatically as the expressive power of the language increases. This leads to a tradeoff between the expressiveness of a representational language and its computational tractability. It is shown that this tradeoff can be seen to underlie the differences among a number of existing representational formalisms, in addition to motivating many of the current research issues in knowledge representation. (37 refs.)

1299 Expert System Models for Inference with Imperfect Knowledge: A Comparative Study
A. Goicoechea, A.P. Sage, D. A. Schum
Sch. of Inf. Techno. & Eng., George Mason U.
Fairfax, Virginia, USA
Proceedings of the 1987 International Conference on Systems, Man, and Cybernetics (Cat. No. 87CH2503-1), Alexandria, VA, USA, 20-23 Oct 1987 (New York, NY, USA: IEEE 1987), pp 559-63 vol.2

A detailed comparative study is presented of six leading methods for reasoning: Bayes' rule, Dempster-Shafer theory, fuzzy set theory, the MYCIN model, Cohen's system of inductive probabilities, and a class of nonmonotonic reasoning methods. Each method is presented and discussed in terms of theoretical content, a detailed numerical example, and a list of strengths and limitations. The same numerical example is addressed by each method, so that the assumptions and computational requirements that are specific to each method can be highlighted. Guidelines are offered to assist in the selection of the method that is most appropriate for a particular problem. (12 refs.)

1337 Knowledge Acquisition by Inductive Learning from Examples
J. Selbig
Dept. of AI, Central Inst. of Cybern. & Inf. Processes
Berlin, East Germany
Analogical and Inductive Inference. International Workshop All '86 Proceedings, Wendisch-Rietz, East Germany, 6-10 Oct 1986 (Berlin, West Germany: Springer-Verlag 1987), pp 145-63

Before the author describes his approach to the problem of learning the action part of IF (pattern) THEN DO (action) rules, he gives a survey of the problem of knowledge acquisition for expert systems. In connection with this work, he focuses on automatic knowledge acquisition by learning methods. (26 refs.)

1509 A Selected Artificial Intelligence Bibliography for Operations Researchers
B. Jaumard
Rutgers U., New Brunswick, New Jersey, USA
Peng Si Ow, B. Simeone
Ann. Oper. Res. (Switzerland), vol. 12, no. 1-4, pp 1-50, 1988

The authors have compiled a selected, classified, and annotated artificial intelligence bibliography specifically addressed to an operations research audience. The bibliography includes approximately 450 references from the areas of search (including heuristics and games), automatic deduction (including theorem proving, logic programming, and logical aspects of databases), planning, learning, and knowledge-based systems (with numerous specific applications to management, engineering, science, medicine, and other fields). They have also added a general references section, as well as a special section on artificial intelligence/operations research interfaces. (457 refs.)

1527 Comments on "A Critique of Pure Reason", and Reply
D. McDermott, J.F. Allen, H.A. Kautz, D.G. Bobrow, M.J. Stefik, K.A. Bowen, R.J. Brachman, E. Charniak, J. de Kleer, J. Doyle, K.D. Forbus, P.J. Hayes, C. Hewitt, G.E. Hinton, J.R. Hobbs, D. Israel, R. Kowalski, J. McCarthy, V. Lifschitz, R.C. Moore, N.J. Nilsson, A. Pentland, D. Pole, R. Reiter, S.J. Rosenschein, L.K. Schubert, B.C. Smith, M.E. Stickel, W.M. Tyson, R. Waldinger, T. Winograd, W.A. Woods
Computational Intelligence (Canada), vol.3, no. 3, pp 161-237, Aug 1987

McDermott (ibid., vol.3, no.3, p.151-60, 1987) argued that attempts to formalize commonsense knowledge using first-order logic had failed and that the logicist approach was based on a false premise. The other authors, except Woods,

defend logicism and criticize McDermott's views in various ways in 26 papers. McDermott's reply is in two parts; a clarification of his earlier arguments, and a discussion of the nature of knowledge representation and whether AI programs differ from any others in this regard. Finally, Woods argues that logic was designed for other types of knowledge representation, and that logicists try in effect to use a good tool for an inappropriate job. He then develops an account of procedural semantics, arguing that it is a better foundation for commonsense knowledge formalism.

1774 A Code of Professionalism for AI
B. Whitby
AISB Q. (UK), no. 64, pp 9-10, Spring 1988

As AI becomes important as a commercial product, a change in attitudes to professionalism may be required. Commercial AI cannot simply adopt the informal, experimental nature of academic AI. Given the need to attempt some more towards professionalism, the publication of a code of conduct seems an easy and effective first step. Obviously to be of value a code should be acceptable to the majority of AI practitioners — a difficult objective, given the diverse nature of the field. The publication of a code of conduct should be a catalyst in the move towards professionalism in AI. It is only after a code is acceptable to those who will actually have to work within it that it will have any real meaning. Before that point it is merely a discussion document. (3 refs.)

1793 OBSERVER: a Probabilistic Learning System for Ordered Events
K.C.C. Chan
Dept. of Syst. Design Eng., U. of Waterloo, Ont., Canada
A.K.C. Wong, D.K.Y. Chiu
Pattern Recognition, 4th International Conference. Proceedings, Cambridge, UK, 28-30 March 1988 (Berlin, West Germany: Springer-Verlag 1988), pp 507-16

Given a sequence of observed events which are ordered with respect to time or position and are described by the coexistence of several discrete-valued attributes that are assumed to be generated by a random process, the inductive prediction problem is to find the probabilistic patterns that characterize the random process, thereby allowing future events to be predicted. This paper presents a probabilistic inference technique for solving such a problem. Based on it, a learning program called OBSERVER has been implemented. OBSERVER can learn, inductively and without supervision, even if some observed events could be erroneous, occasionally missing, or subject to certain degrees of uncertainty. It is able to reveal the patterns and regularities inherent in a sequence of observed events and cannot only specify, in a clearly defined way, the happenings in the past but also gain insight for prediction. The proposed technique can be applied to solve different problems in artificial intelligence and pattern recognition where decisions concerning the future have to be made.

2.0 SYSTEMS AND TECHNIQUES

1358 Systems Analysis Techniques for the Implementation of Expert Systems
D. R. Martinez
Arco Oil & Gas Co., Dallas, Texas, USA
M. G. Sobol
Inf. Softw. Technol. (UK), vol.30, no. 2, pp 81-8, March 1988

Knowledge-based expert systems incorporate human expert knowledge with the use of computer systems. Today, the wide availability of expert systems shells allows the knowledge engineer to implement specific rules for a desired application. The final design of the artificial intelligence system is the outcome of a detailed study in an organized fashion. The application of system analysis tools for designing knowledge-based expert systems is presented. The paper illustrates the application of the tools with a simplified example drawn from the oil and gas exploration business. The use of a systematic approach in designing expert systems should help the knowledge engineer clearly identify the facts and rules representative of the acquired human knowledge. (15 refs.)

1360 Trends in Artifical Intelligence
I/S Anal. (USA), vol. 26, no. 2, pp 1-12, Feb 1988

Until recently, AI applications were generally thought to require special hardware and software and work best in stand-alone applications. Recently, however, that belief has been reversed. Now, many of the benefits of AI are seen to come from using it to extend and enhance conventional information systems. The report looks at two areas in AI that are likely to become integral parts of many information systems — expert systems and natural language processing systems. In addition, tools used to develop expert systems may foretell features of future, all-purpose programming environments. (14 refs.)

1379 Management of Uncertainty in a Medical Expert System
D.L. Hudson, M.E. Cohen
California U., San Francisco, California, USA
International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems, Paris, France,30 June-4 July 1986, Berlin, West Germany: Springer-Verlag 1987, pp 283-93

The use of uncertainty in a rule-based expert system for the analysis of chest pain is discussed. The system, EMERGE, has been evaluated retrospectively and prospectively and has been found to perform extremely well. The original system has been altered to handle degrees of presence of symptoms and variable contribution of antecedents. It also utilizes a logical construct which generalizes traditional AND/OR logic. (26 refs.)

1391 Toward Considering Psychological Measurement Issues when Developing Expert Systems
L.Adelman
Dept. of Inf. Syst. & Syst. Eng., George Mason U. Fairfax, Virginia, USA
Proceedings of the 1987 International Conference on Systems, Man, and Cybernetics (Cat. No. 87CH2503-1), Alexandria, VA, USA, 20-23 Oct 1987 (New York, NY, USA: IEEE 1987), pp 1044-8 Vol.3

There are four possible sources (or determinants) of the quality of different components of the knowledge base contained in an expert system: the domain experts, the knowledge engineers, the methods used by the knowledge engineers when working with the domain experts, and the problem setting and structure used to facilitate knowledge elicitation. A study has been undertaken to assess the extent to which system validity is affected by variability of any of one of these. Toward this end, the results of reanalyzing the data from an experiment varying domain experts, knowledge engineers, and elicitation methods when developing multi-attributed hierarchies is presented. No significant effects were obtained for elicitation method or knowledge engineer, and the results suggest that domain experts were the largest source of variation. (12 refs.)

1415 Experiences with OPS5 in the Development of Expert Systems at Bayer
F. Biegler-Konig
Bayer AG, Leverkusen, West Germany
Expert Systems in Production Engineering. Proceedings of the International Workshop, Spa, Belgium, 18-22 Aug 1986, (Berlin, West Germany: Springer-Verlag 1987), pp 157-61

Techniques of AI are supposed to be one of the most important technologies of the future. Great interest has been shown in the use of expert systems, but much effort must still be invested to get more experience in their development and applications. At

Bayer some studies were made about the advantages of the application of expert systems in chemical engineering. Many areas of possible applications were located, especially for small, compact and easy to use systems with connections to conventional software (Fortran, graphics). Small systems only need a limited amount of development effort and are therefore better suited to gain experiences and assess the power of the new technology. (4 refs.)

1549 A Company/University Joint Venture to Build a Knowledge-Based System
J.R. Weitzel, K.R. Andrews
South Carolina U., Columbia, South Carolina, USA
Manage. Inf. Syst. Q. (USA), vol. 12, no. 1, pp 23-34, March 1988

A joint venture between a university-based research institute and a health insurance company to build a knowledge-based system to perform medical review of health insurance claims is described. The article examines the impact of the differing cultures of the company and the university on the cohesion in the joint knowledge engineering group. It also examines the current literature on the development cycle for building knowledge-based systems as the framework for analyzing the events in the project, particularly the influence of the claims review task on system design. From another perspective, it examines participant roles in terms of shifts of attention among domain knowledge, knowledge representation, system performance, and the kinds of skills needed to improve the evolving system. The conclusion includes a series of recommendations that may assist other companies and universities setting up similar joint ventures. (16 refs.)

1553 Elements of Expert System Shells
M. Fontana, J. Zeimetz
Periscope Co. Inc., Atlanta, Georgia, USA
PC Tech J. (USA), vol. 6, no. 5, pp 63-5, May 1988

To build the appropriate expert system for a given problem, the developer must first select the proper tool, known as an expert system shell. Deciding which shell is best means knowing which questions to ask. To help ask the right questions, the authors present criteria for evaluating shells. In particular, they discuss knowledge representation methods, inference engines and tracing.

1557 The KREME Knowledge Editing Environment
G. Abrett, M.H. Burstein
BBN Labs., Cambridge, Massachusetts, USA
Int. J. Man-Mach. Stud. (UK), vol. 27, no.2, pp 103-26, Aug 1987). *AAAI Workshop on Knowledge Acquisition for Knowledge-Based Systems*, Banff, Alberta, Canada, Nov 1986

One of the major bottlenecks in large-scale expert-system development is the problem of knowledge acquisition: the construction, maintenance, and testing of large knowledge bases. This paper provides an overview of the current state of development of the Knowledge Representation Editing and Modeling Environment (KREME). It is an extensible experimental environment for developing and editing large knowledge bases in a variety of representation styles. It provides tools for effective viewing and browsing in each kind of representational base, automatic consistency checking, macro-editing facilities to reduce the burden of large-scale knowledge-base revision and some experimental automatic generalization and acquisition facilities. (22 refs.)

1571 Making Friends of Man and Machine [Expert Systems]
S. Ottley
Syst. Int. (UK), vol. 16, no. 4., pp 23-5, April 1988

The author warns of some human and organizational problems to consider when implementing an expert system, from capturing an expert's knowledge to gaining user acceptance.

1577 Expert Systems Maintenance: Post-implementation Support
J.E. Caviedes
North American Philips Corp.
Briarcliff Manor, New York, USA
IEEE MONTECH '87 Conferences: COMPINT (Cat. No. 87CH2518-9), Montréal, Que., Canada, 9-12 Nov 1987 (New York, NY, USA: IEEE 1987), pp 21-4

Knowledge-base maintenance is a growing concern within the community of expert systems developers. One of the factors limiting the initial rush to apply expert systems technology is the difficulty of post-delivery maintenance. Since the quality of an expert system depends more on the knowledge engineering skills of its developers than on the implementation toolset, it is almost impossible to predict its maintainability. It is argued that maintenance tools can be developed in parallel with the expert system if it has a mature and stable architecture. (10 refs.)

1578 The Software Engineering of Systems with Expert Components
W.B. Frekes, C.J. Fox
AT&T Bell Labs., Holmdel, New Jersey, USA
Proceedings of the Twenty-First Annual Hawaii International Conference on System Sciences. Vol. II. Software Track (Cat. No. 88TH0212-1, Kailu-Kona, HI, USA, 5-8 Jan 1988 (Washington, DC, USA: IEEE Comput. Soc. Press 1988), pp 48-53

Integrating expert system components into production software can be difficult because expert system development environment are typically incompatible with traditional software engineering technology. In an effort to deal with this problem, the authors are developing CEST, a C expert system toolset. CEST is a library of inference engines implemented as C functions callable from C programs, and a workbench of knowledge engineering support tools. CEST allows easy integration of expert system components into C-based software systems and provides knowledge-engineering support tools analogous to traditional software-engineering support tools. (27 refs.)

1584 A Comparison of the Manipulation of Certainty Factors by Individuals and Expert System Shells
D. Kopcso, L. Pipino
Babson Coll., Babson Park, Massachusetts, USA
W. Rybolt
Proceedings of the Twenty-First Annual Hawaii International Conference on System Sciences. Vol. III. Decision Support and Knowledge Based Systems Track (Cat. No. iiTH0213 9), Kailua-Kona, HI, USA, 5-8 Jan 1988 (Washington, DC, USA: IEEE Comput. Soc. Press 1988), pp 181-8

The treatment of uncertainty in expert system shells is addressed, starting with a review of the modeling of uncertainty by expert system shells. An experiment to replicate earlier work investigating the manner in which individuals manipulate certainty factors in comparison to commercial shells is discussed. Comparisons are made among seven commercial shells, both personal-computer (PC)-based and mainframe-based, and individuals. A significant difference between individuals and shells themselves is indicated. Some implications for both expert system and decision-support-system methodologies are discussed. (25 refs.)

1829 Knowledge Harvesting: a Practical Guide to Interviewing
M. Davies
Software A & E, Chichester, UK
S. Hakiel
Expert Syst. (UK), vol. 4, no.1, pp 42-50, Feb 1988

It is up to the knowledge engineer not only to do his best to get

the relevant information from the domain expert, but also to make the expert's part of the interaction as pleasant and as easy as possible, and to generate and maintain in him as much interest and enthusiasm as possible. The aim of this article is to provide some guidelines on how to do all this, culled from the experience of the author and the reported views of other workers in the area. Because this is intended as a practical guide rather than a learned treatise, it is orientated primarily towards quick reference. Each section covers one major part of the knowledge harvesting task; within each section, the text is principally organized into notes. It is recommended that the reader make a quick pass through the whole document before starting a knowledge harvesting exercise, and refer to specific parts in more detail as and when required during the course of the exercise. This report does not go into the absolute merits of various approaches. (3 refs.)

1851 Artificial Intelligence: Current Technologies and Tools
W.E. Bracker, Jr.; L.C. Bracker
Tech. Res. Associates, Tucson, Arizona, USA
B.R. Konsynski
Proceedings of the Seventh Annual Conference and Exposition: Computer Graphics'86. Anaheim, CA, USA, 11-15 May 1986 (Fairfax, VA, USA; Nat. Comput. Graphics Asoc. 1986), pp 37-47, vol.3

AI has many subfields. Among them are: computer vision, voice-recognition, natural language, robotics and expert systems. The authors provide an introduction to expert systems, discuss languages used in expert system development, consider current technologies in support of AI, and present via a comparison matrix, AI workstations. (83 refs.)

1859 SEA — an Expert System for Nuclear Test Ban Treaty Verification
C.L. Mason, R.R. Johnson
Dept. of Appl. Sci., California U.
Livermore, California, USA
R.M. Searfus, D. Lager
Proceedings of the Australia Joint Artificial Intelligence Conference — AI '87, Sydney, NSW, Australia, 2-4 Nov 1987 (Sydney, NSW, Australia: U. of Sydney 1987), pp 11-25

Presents an expert system that interprets seismic data from Norway's regional seismic array, NORESS, for underground nuclear weapons test ban treaty verification. Three important aspects of the expert system are that it emulates the problem solving behaviour of the human seismic analyst, it acts as an assistant to the human analyst by automatically interpreting and presenting events for review, and it enables the analyst to interactively query the system's chain of reasoning and manually perform an interpretation. The general problem of seismic interpretation is described. The expert system is presented in terms of problem solving strategy, representation structures and user interface elements. (12 refs.)

1868 Induction, Knowledge and Expert Systems
J.R. Quinian
Key Centre for Advanced Comput. Sci.
New South Wales Institute of Technology
Sydney, New South Wales, Australia
Proceedings of the Australian Joint Artificial Intelligence Conference — AI'87, Sydney, NSW, Australia, 2-4 Nov 1987 (Sydney, NSW, Australia:U. of Sydney 1987), pp 223-40

Expert systems technology is a cornerstone of applied artificial intelligence that has many potential benefits for the manufacturing and service industries. The principal barrier to the more widespread use of expert systems is the difficulty of assembling and debugging the requisite collections of expert knowledge. Induction, the process of generalizing from examples, provides a proven means of expediting the knowledge acquisition process. this paper sketches the key ideas, supported by a case study, and discusses experimental techniques for improving the intelligibility of the induced knowledge. (29 refs.)

3.0 APPLICATIONS

1419 REGWASTE: An Expert System for Regulating Hazardous Wastes
G. Anandalingam
Dept. of Syst., Pennsylvania U.
Philadelphia, Pennsylvania, USA
Proceedings of the 1987 International Conference on Systems, Man, and Cybernetics (Cat. No. 87CH2503-1), Alexandria, VA, USA, 20-23 Oct 1987 (New York, NY, USA: IEEE 1987), pp 634-9, vol. 2

The author reports on an expert system called REGWASTE, which has been implemented to assist the US Environmental Protection Agency in regulating hazardous waste. The expert system simplifies a complex decision problem, provides alternatives arranged according to likelihood, and advises the user on further information needed to finalize the regulation. (6 refs.)

1421 A "Neural" Network that Learns to Play Backgammon
G. Tesauro
Center for Complex Syst. Res., Illinois U.
Champaign, Illinois, USA
T. J. Sejnowski
1987 IEEE Conference on Neural Information Processing Systems — Natural and Synthetic Abstracts of Papers (Cat. No. 87CH2386-1), Denver, CO, USA, 8-12 Nov 1987 (New York, NY, USA: IEEE 1987), pp 28

Summary form only given. A class of connectionist networks that have learned to play backgammon at an intermediate-to-advanced level is described. The networks were trained by a supervised learning procedure on a large set of sample positions evaluated by a human expert. In actual match play against humans and conventional computer programs, the networks demonstrate substantial ability to generalize on the basis of expert knowledge. The study touches on some of the most important issues in network learning theory, including the development of efficient coding schemes and training procedures, scaling, generalization, the use of real-valued inputs and outputs, and techniques for escaping from local minima. Practical applications in games and other domains are also discussed.

1425 Artificial Intelligence: Current Trends Expanding
J. K. Carter
Office (USA), vol. 107, no. 1, pp 82, Jan 1988

The field of AI has experienced its share of exciting trends and developments with research in areas such as neural networks. Of major importance to outside research labs is progress towards integration of AI-based system functions into conventional data processing and office system environments. The trend toward integration of AI and conventional information processing is allowing computerized business systems to expand beyond clerical tasks into more knowledge-intensive applications. Because of their ability to support a decision-making process from inception to conclusion, expert systems, integrated with conventional data processing systems add significant value to a business by leveraging the valuable time of professional personnel.

1448 Applying AI Techniques for Patent Information Retrieval
D. Vermeir, E. Laenens, J. Dierick
Dept. of Math. & Comput. Sci., Antwerp U.
Wilrijk, Belgium
World Pat. Inf. (USA), vol. 10, no.1, pp 26-36, 1988

Presents an overview of recent developments in software technology, especially information retrieval and expert

systems. Particular consideration is given to the possible applications in the area of user-friendly access to patent information systems. A proposal for an expert system, that could act as a knowledgeable intermediary between the end user with no information retrieval experience and the various host systems, is described. (16 refs.)

1449 Overview of Artificial Intelligence and Expert Systems for Information Professionals
D.T. Hawkins
AT&S Bell Labs, Murray Hill, New Jersey, USA
Online '87 Conference Proceedings,
Anaheim, CA, USA, 20-22 Oct 1987
(Weston, CT, USA: Online 1987), pp 94-6

Interest in AI and expert systems is growing rapidly as AI-based systems move out of the research laboratory and into the marketplace. AI will be a major research effort in the coming years, with information retrieval a prime application. As interest in AI and its related areas grows, online searchers and other information professionals need to understand the terminology and basic principles so they can intelligently evaluate the new systems that are appearing. The author introduces AI as it relates to information retrieval and online systems. It defines and explains the terms commonly used in AI, describes some AI programming languages and hardware, and briefly discusses some well known prototype systems. Applications of AI for information retrieval are then described, and finally, some commercial software packages now beginning to appear on the market are introduced. (5 refs.)

1457 A Prototype Expert System for Synthesizing Chemical Process Flowsheets
R. L. Kirkwood, M.H.Lockhe, J.M. Douglas
Dept. of Chem, Eng., Massachusetts U.
Amherst, Massachusetts, USA
Comput. Chem. Eng. (UK), vol. 12, no. 4, pp 329-43, Apr 1988

PIP (process invention procedure) is an hierarchical expert system for the synthesis of chemical process flowsheets. It uses a combination of qualitative and quantitative knowledge, arranged in a hierarchical structure. The heuristics are used to select the unit operations, to identify the interconnections between these units, to identify the dominant design variables are to identify the process alternatives at each level of the hierarchy, while the quantitative models are used to calculate process flows, the equipment sizes and costs the raw material and utility costs and the process profitability at each level. A hybrid, expert system control architecture was developed for PIP that allows these two types of knowledge-bases to interact in such a way that the heuristic rules "fire" the appropriate subroutines used to evaluate the flowsheet. PIP attempts to invent a flowsheet using a depth-first strategy, where one of the goals is to see if there is some reason why none of the alternatives will ever be profitable. Thus, it attempts to complete a design before any alternatives are considered. If profitable operation is observed over some range of the design variables at a particular level, the PIP proceeds to the next level and adds more detail to the flowsheet. However, if the process is not profitable, PIP examines the alternatives, starting with the ones found at the earliest level, and if no profitable alternative can be found the design project is terminated. Hence, PIP allows a design engineer to invest initial chemical flowsheet structures rapidly, to estimate the optimum design conditions for this flowsheet, to identify possible process alternatives and to quickly screen these alternatives. (33 refs.)

1458 EMPS — Making Automated NC Programming Feasible
FMS Mag. (UK), vol. 5, no.1, p.25-6, Jan 1988

Automation of the NC programming function is the goal of a new project initiated by Computer Aided Manufacturing-International. The system will be called EMPS (expert manufacturing programming system). So far, the project has developed a comprehensive design for an automated processor. Using techniques such as solid modeling geometric tolerancing, expert systems, features recognition, and volume decomposition, the processor will automatically determine the volumes of material to be removed, the appropriate NC machine fixturing requirements, cutting tool assemblies, cutter paths, and machinability data (feed and speed rates). The challenges lie in three areas: task decomposition, work element generation and tool path generation.

1490 An Approach to Integrating Expert System Components into Production Software
W. B. Frakes, C.J. Fox
AT&T Bell Labs., Holmdel, New Jersey, USA
Proceedings 1987 Fall Joint Computer Conference — Exploring Technology: Today and Tomorrow
(Cat. No. 87CH2468-7), Dallas, TX, USA, 25-29 Oct 1987
(Washington, DC, USA: IEEE Comput. Soc. Press 1987), pp 50-6

The authors describe work on CEST (C expert system tools), an expert system function library and workbench for the Unix/C environment that has been developed to incorporate expert systems techniques into production software. CEST is being designed to consist of a library of inference engines implemented as C functions and callable from C programs, and a workbench of knowledge engineering support tools for building, analyzing, and maintaining knowledge bases. The authors describe the first tool in CEST, called AVIEN, which is a backward chaining inference engine implemented as a set of C functions. (21 refs.)

1474 Rocket Engine Health Monitoring System (HMS) via an Embedded Expert System (EES)
J. Pooley, T. Homsley, W. Tech
SPARTA Inc., Huntsville, Alabama, USA
J. Jones, P. Lewallen
Proceedings of the 1987 International Conference on Systems, Man. and Cybernetics (Cat. No. 87CH2503-1),
Alexandria, VA, USA, 20-23 Oct 1987 (New York, NY,
USA: IEEE 1987), pp 1127-32 vol. 3

The SPARTA embedded expert system (SEES) is an intelligent system that directs the analysis of rocket engine maintenance requirements by placing confidence factors on possible engine status and then recommending a course of action to an engineer or engine controller. This technique can prevent catastrophic failures or costly rocket engine downtime because of false alarms. Further, the SEES has potential as an onboard flight monitor for reusable rocket engine systems. The SEES methodology synergistically integrates vibration analysis, pattern recognition, and communications theory techniques with an artificial intelligence technique. Results from processing real data are presented. (5 refs.)

1465 Air Traffic Control using AI Techniques
M. H. Davies
Sist. Autom. (Italy), vol. 33, no. 287, pp 1292-302,
Dec 1987, in Italian

Describes a project carried out by Systems Designers Ltd. in collaboration with RSRE (Royal Signals and Radar Establishment) which has already resulted in a prototype air traffic control system using AI techniques. The process of command and control is considered and the requirements of the air traffic control operation described in detail. An expert system type aid, ATC Aid, is proposed, with a description of its architecture, mode of operation, and the way it interacts with the human controller. Verification and validation of the system are considered. It is concluded that the feasibility of an AI approach to air traffic control has been demonstrated. (10 refs.)

1500 Extensions and Modifications to the Prolog Language in View of Actual Application Programming
Expert Systems in Production Engineering. Proceedings of the International Workshop, Spa, Belgium, 18-22 Aug 1986 (Berlin, West Germany: Springer-Verlag 1987), pp 145-56

The initial work that the authors have done has proved that expert system technology should be regarded as an extension of conventional data processing. With these techniques it has become possible to solve problems which could not be solved before by conventional methods. As with all new technologies, the techniques of AI and especially expert systems are based on an initial theory developed over recent years. This theory was mainly based on a general view of the possible thematics derived by mathematicians and informatic engineers. In view of the recent work carried out by the authors, they find that a number of extensions and modifications have to be executed to make the use of the Prolog language more efficient in the hands of the users.

1609 Expert Systems and the Law
A. Bloch
Micro Syst. (France), no. 84, pp 203-4, March 1988, In French

Legal ramifications of the use of expert systems are considered with particular reference to the two sections: the knowledge base, and the inferential mechanism. Contractual details regarding construction of specific expert systems are being developed with the emphasis on turnkey contracts.

1612 The Alvey DHSS Demonstrator Project: Applying Intelligent Knowledge-based Systems to Social Security
G.N. Gilbert
Dept. of Sociology, Surrey U., Guildford, UK
Inf. Age (UK), vol. 10, no. 2, pp 113-15, April 1988

Prototype decision support systems being developed under the Alvey Program are outlined. They are designed for possible use by the UK Department of Health and Social Security and cover the assessment of benefit claims in local offices, the formulation of legislation by policy-makers, and the provision of advice and assistance to claimants. An attempt is being made to utilize the strengths of both user and system in the design. (6 refs.)

1618 QUARTZ: an Intelligent Assistant for the Analysis and Evaluation of Project Proposals
K. Dalkir, J. Muzard
Canadian Workplace Autom. Res. Centre, Laval, Que.
IEEE MONTECH '87 Conferences: COMPINT (Cat. No. 87CH2518-9), Montréal, Que., Canada, 9-12 Nov 1987 (New York, NY, USA: IEEE 1987), pp 208-10

QUARTZ is an intelligent decision-support system that will assist Canadian government analysis in their evaluation of project proposals submitted for funding under the economic regional development assistance program for Quebec. The authors primarily focus on the knowledge engineering methodology used to develop QUARTZ. In particular, the knowledge acquisition, knowledge representation, and knowledge validation stages of system development are discussed, both in terms of lessons learned and future research directions to be undertaken. (7 refs.)

1621 Potential Defense Applications of Expert Systems
V. Shah
Southwest Texas State U., San Marcos, Texas, USA
G.D. Buckner
IEEE Aerosp. Electron. Syst. Mag. (USA), vol.3, no.2, pp 15-21, Feb 1988

The authors provide an overview of expert systems and how they may effect the development of future defense applications. Military uses of computers are outlined, and expert systems fundamentals are described. AI research and development efforts by the military are examined, and potential military applications are discussed. Expert systems efforts at NASA, by the US Air Force, and for the Strategic Defence Initiative are considered. (19 refs.)

1649 A Financial Investment Assistant
K. Kandt, P. Yuenger
Teknowledge Fed. Syst., Thousands Oaks, California, USA
Proceedings of the Twenty-First Annual Hawaii International Conference on System Sciences. Vol. III. Decision Support and knowledge Based Systems Track (Cat. No.88TH0213 9), Kailua-Kona, HI, USA, 5-8 Jan 1988 (Washington, DC, USA: IEEE Comput. Soc. Press 1988), pp 510-17

The analysis of financial markets is a time-consuming complex and error-prone process. The system described which is still under development, is an attempt to improve this process by partially automating the acquisition, analysis, and selection of financial market investments. The ultimate goal is to fully automate this activity. The current approach is to use technical analysis and fundamental analysis to determine when to buy, sell, or hold various instruments, and AI techniques to select a portfolio of stocks and/or options based on the goals of the system user. The tool uses a dynamic interface that is reactive to human interaction. (25 refs.)

1671 Expert Systems for University Admissions
J.S. Edwards, J.L. Bader
Aston U., Birmingham, UK
J. Oper. Res. Soc. (UK), vol. 39, no. 1, pp 33-40, Jan 1988

Describes the construction of an expert system to help the admissions tutor for a university degree in business and management which receives some 2000 applications for entry each year, using the SAGE shell. What originally began as a "demonstrator project" is shown to be of practical value, in terms of both producing a usable expert system and clarifying and questioning the selection criteria used by the admissions tutor. A particular conclusion emerging from this work which may be relevant to many expert systems applications is that the ethical and practical considerations dictated that some questions involving judgement could not be delegated by the admissions tutor to the clerical staff. It thus became necessary to develop two versions of the expert system, one a full "administrations-tutor system", the other a more limited version for day-to-day use. (9 refs.)

1701 Expert System Tools Emerge from the Technology of AI
F.J. Bartos
Control Eng. (USA), vol. 34, no. 14, pp 36-8, Dec 1987

AI is a technology that works to make computers behave more like humans in executing and solving problems. AI is fundamentally different from conventional programming. It is a more symbolic than numeric process, employing heuristic or rule of thumb procedures over explicit ones to arrive at solutions. It is meant to be easily modified or expanded and can live with less than optimum results — even accepting an incorrect answer or two. It is emphasized that AI, in the form of its expert systems, is starting to go to work for process and control engineers.

1705 Expert System for Wire Cutting EDM, Based on Pulse Classification and Thermal Modeling
W. Dekeyser, R. Snoeys, M. Jennes
Inst. voor Werktuigkunde, Heverlee, Belgium
Robot. Comput.-Integr. Manuf. (UK), vol. 4, no. 1-2, pp 219-24, 1988, (Manufacturing Science, Technology and Systems of the Future, Ljubljana, Yugoslavia, 12-14 Sept 1985)

Although wire electric discharge machine (EDM) has become a fully competitive machining technique during the last decade, the process performance is limited by some typical restrictions. In order to improve machining performance, a multi-disciplinary expert system was developed. It mainly consists of two parts: an observation procedure and a decision procedure. The observation of the wire EDM process is carried

out by means of a powerful electronic device, called the EDM pulse discriminating system. This device delivers "on-line" a "chaos" of information of the process. In the decision phase, the incoming data are analyzed and patterns are detected by means of a computer and eventually the parameter settings of the EDM machine are changed in order to obtain more optimal working conditions. This decision procedure is partially based on the results of a mathematical thermal model of the wire, enabling one to predict the influence of some parameter changes. Further, a technological database is permanently accessible for the control strategy program. Machining speed can be increased without enhancing the risk of wire rupture, because in the case of an increased danger of rupture, appropriate action is taken by the control strategy. The paper gives some results on how the thermal model may predict some thermal overload of the wire. It also describes how frequently observed breakage at sharp corners of the wire path may be dealt with. It further illustrates the actions taken by the expert system when various types of disturbance occur. The consideration of practical experience with a thermal model yields an increased level of machine autonomy ensuring a performance level obtainable only with a skilled operator. (7 refs.)

1710 Glass Annealing Process Simulation Using Expert System: a Glass Industry Application of Artificial Intelligence
R.A. Herrod, J.W. Rickel
Texas Instrum. Inc., Dallas, Texas, USA
T. Garland
IEEE Trans. Ind. Appl. (USA), vol. 24, no. 1, pt. 1, pp 43-8, Jan-Feb 1988

How an AI system was used for simulation of a glass annealing process is discussed. The system consists of two parts: a planner that uses rules-of-thumb (in the form of an expert system) to determine control settings for the glass processing, and a simulator that uses control settings to derive a glass temperature curve. The two subsystems work together to produce the necessary Lehr control settings to anneal the glass product. A short review of expert system technology is given, followed by a discussion of the project and lessons learned. (4 refs.)

1726 RATIONALE: Developing Expert Systems that Reason by Explaining
S. Abu-Hakima
Dept. of Syst. & Comput. Eng., Carleton U., Ottawa, Ontario
F. Oppacher
IEEE MONTECH '87 Conferences: COMPINT (Cat. No. 87CH2518-9) Montreal, Que., Canada, 9-12 Nov 1987 (New York, NY, USA: IEEE 1987), pp 13-16

Explanations generated by current expert systems are often terse, include bookkeeping information irrelevant to most users, and are isolated from the reasoning processes of the expert system. A methodology is presented for building knowledge-based systems that reason by explaining. The methodology is the basis for RATIONALE, a tool that is implemented in Prolog and ties together the processes of knowledge acquisition, expert system construction, and explanation of system reasoning. RATIONALE does this by using domain knowledge together with explicitly represented strategies. Reasoning proceeds by constructing a hypothesis tree having a root that contains the most general diagnosis of the expert system. Guided by a focusing algorithm, the tree branches into more specific hypotheses that explain the more detailed symptoms provided by the user. As the expert system is used, the hypothesis tree also serves as a dynamically generated explanation tree. (22 refs)

1733 A Knowledge-based Approach for Real-time Systems Debugging
J.P. Tsai, K.Y. Fang, V.R.K. Thalla, H. Gandhi
Dept. of Electr. Eng. & Comput. Sci., Illinois U. Chicago, Illinois, USA
Proceedings of the Twenty-First Annual Hawaii International Conference on System Sciences. Vol. II. Software Track (Cat. No. iiTH0212-1), Kailua-Kona, HI, USA, 5-8 Jan 1988 (Washington, DC, USA: IEEE Comput. Soc. Press 1988), pp 533-40

The authors consider a method of testing unpredictable sequences called program execution monitoring. In this method, a snapshot of the system events and happenings are recorded in a non-interfering manner, i.e. without corrupting the critical timing requirements of the system. The traces thus collected are analyzed to isolate the bug in an off-line mode. The examination of voluminous traces is a laborious and tedious task that requires a high degree of expertise. It is shown how this expertise can be encapsulated in a knowledge-based system and the examination process automated to at least localize the fault and answer questions about its reasoning.

1885 Victorian Business Assistance Referral System: Development and Early Use
M. Frazer
Dept. of Ind., Technol. & Resources
Melbourne, Victoria, Australia
Proceedings of the Australian Joint Artificial Intelligence Conference — AI '87, Sydney, NSW, Australia, 2-4 Nov 1987 (Sydney, NSW, Australia: U. of Sydney 1987), pp 44-60

An expert system to provide advice on assistance to new or small businesses in use at a number of sites in Melbourne. The system has been developed by BBJ Computers International Pty. Ltd of South Melbourne, for the Victorian Government's Department of Industry Technology and Resources. The paper discusses three matters: the Department's motivation in commissioning the system; the planning and development of the system; and early experience of its use. The presentation is from the point of view of the Department as the client, rather than from the point of view of the technical developers of the system. (3 refs.)

1887 Expert Systems for Business Applications
J. Liebowitz
Dept. of Manage. Sci., George Washington U. Washington, DC, USA
Appl. Artif. Intell. (USA), vol. 1, no. 4, pp 307-13, 1987

Expert systems are becoming more prevalent in financial and business applications. With venture capital pouring into the expert systems area, there is greater interest by commercial firms. Companies are either building in-house expert systems capabilities or contracting out to expert system developers to construct expert systems for business applications. This paper presents much of the work being done in developing expert systems in business and discusses fertile areas for constructing more business expert systems. (10 refs.)

1897 Expert Process Planning System with Solid Model Interface
S. Joshi, N.N. Vissa, Tien-Chien Chang
Sch. of Ind. Eng., Purdue U., W. Lafayette, Indiana, USA
Int. J. Prod. Res. (UK), vol. 26, no. 5, pp 863-85, May 1988

Presents an integrated hierarchical framework of a process planning system with a CAD interface. The objective of the project discussed is to integrate design with process planning using AI techniques. The development of a CAD interface is discussed with respect to automated feature recognition,

determination of tool approach direction, and deciding the precedence relationship between the features. Sample results from the CAD interface are presented. The expert system for the process planning module is discussed with the part representation and knowledge base, and the plan generation procedure. The module uses hierarchically organized frames for both part representation and the knowledge base. (40 refs.)

1899 Knowledge Based Simulation Techniques for Manufacturing
R.E. Shannon
Dept. of Ind. Eng., Texas A&M U.
College Station, Texas, USA
Int. J. Prod. Res. (UK), vol. 26, no.5, pp 953-73, May 1988

The art and science of simulating complex manufacturing systems is rapidly changing. A great deal of attention is being devoted to the possibilities of bringing AI and expert systems technology into simulation methodology. Such systems will hopefully allow models to be quickly developed, validated and run with as much of the necessary expertise as possible built into the software. This paper addresses: the motivation and need for developing such systems; the nature of such systems; the potential benefits of this technology over existing approaches; and the current state-of-the-art as it applies to simulation. (60 refs.)

1909 The Expert System for Career Planning: ACES
S. Geffin, T. Burges, B. Furht
Dept. of Electr. & Comput. Eng., Miami U.
Coral Gables, Florida, USA
Microcomput. Appl. (USA), vol. 6, no.3, pp 71-7, 1987

The design of a knowledge-based student advisor system implemented on a microcomputer is discussed. The student advisor system is intended to assist graduating students, majoring in either electrical, computer engineering, or computer science in deciding upon an area of specialization. The system can also prescribe a list of graduate schools which match the student's qualifications, location desired and the area of specialization decided by the system. The programming tool used to develop the system was the Texas Instrument Personal Consultant run on an IBM/AT microcomputer. (11 refs.)

1926 GOLD: an Expert System for Mineral Identification from Reflectance Spectra
R. Lister
Basser Dept. of Comp. Sci., Sydney U., NSW, Australia
K. Ali, R. Buda, C. Horsfall, W. Buntine
Proceedings of the Australian Joint Artificial Intelligence Conference — AI '87, Sydney, NSW, Australia), 2-4 Nov 1987
(Sydney, NSW, Australia: U. of Sydney 1987), pp 29-43

GOLD is an expert system for identifying minerals, some of which are indicative of fine-grained sub-surface gold deposits. The expert system identifies minerals by searching for characteristic features in the infrared reflectance spectra of rock samples. The system is being built using the Prospector expert system shell which represents knowledge using an inference net and propagates uncertainty using subjective Bayesian reasoning. A prototype has been installed at CSIRO and is currently on trial. GOLD's performance compares favourably with that of domain experts. The system is being extended to deal with spectra produced by samples containing more than one mineral. (11 refs.)

1932 Application of Artificial Intelligence in Telecommunications
K.J. Macleish, S. Thiedke, D. Vennergrund
GTE Commun. Syst., Phoenix, Arizona, USA
GLOBECOM Tokyo '87. IEEE/IECE Global Telecommunications conference 1987. Conference Record
(Cat. No. 87CH2520-5), Tokyo, Japan, 15-18 Nov 1987
(New York, NY, USA: IEEE 1987), pp 329-33, vol. 1

The authors examine the ways in which AI techniques are used now and will be used in the future to assist in many aspects of telecommunications systems. The current use of AI consists primarily of expert systems designed to aid in diagnosing complex equipment. The authors briefly examine such a system, NEMESYS, which diagnoses intermittent faults in the GTD-5 EAX digital central office switch, and focus on why this project required AI techniques. Furthermore, the authors explore how other AI techniques, such as natural language processing, distributed AI, and speech recognition, will be used along with expert systems to assist both operations and services in what they call the network of the future. (10 refs.)

1939 Knowledge Based Maintenance in Networks
D. Peacocke, S. Rabie
Bell-Northern Res., Ottawa, Ontario, Canada
GLOBECOM Tokyo '87 IEEE/IECE Global Telecommunications Conference 1987. Conference Record
(Cat. No. 87CH2520-5), Tokyo, Japan, 15-18 Nov 1987
(New York, NY, USA: IEEE 1987), pp 1833-8, vol. 3

The maintenance advisor for DMS-100 is considered. MAD is an interactive expert system for helping operating company personnel perform maintenance on the DMS-100 family of digital switches. A prototype is described that is based on using product documentation and experience obtained from craftspeople and field technical support to create the knowledge base for an expert system that diagnoses and chooses corrective actions for DMS-100 problems. Its operation is described, and steps for further evolution of the system in a network maintenance role are outlined. (4 refs.)

1947 ESA: Expert Structural Analysis for Engineers
B.W.R. Forde, S.F. Stiemer
Dept. of Civil Eng., UBC, Vancouver, British Columbia
Comput. Struct. (UK), vol. 29, no. 1, pp 171-4, 1988

The finite element method of analysis is possibly the most popular tool available for the numerical solution of complex problems in engineering. A consistent algorithmic approach to formulation, discretization, and solution procedures for this method has resulted in the wide-spread development of extremely powerful computer software. This power has its price: expertise is required for efficient analysis. A knowledge-based expert system called expert structural analysis (ESA) is being developed which will allow average structural engineers to skillfully use a nonlinear finite element analysis program. The paper examines the nature of the algorithms employed in the analysis program and introduces the principle issues associated with human and artificial expertise. A case study is used to demonstrate the application of a prototype expert system. (9 refs.)

1949 Teaching Artificial Neural Systems to Drive: Manual Training Techniques for Autonomous Systems
J.F. Shepanski, S.A. Macy
TRW Inc., Redondo Beach, California, USA
vol.848, pp 286-93, 1988, *Intelligent Robots and Computer Vision*. Sixth in a Series, Cambridge, MA, USA, 2-6 Nov 1987

The authors have developed a methodology for manually training autonomous control systems based on artificial neural systems (ANS). In applications where the rule set governing an expert's decisions is difficult to formulate, ANS can be used to extract rules by associating the information as expert receives with the actions he takes. Properly constructed networks imitate rules of behaviour that permit them to function autonomously when they are trained on the spanning set of possible situations. This training can be provided manually, either under the direct supervision of a system trainer, or indirectly using a background mode where the network assimilates training data as the expert forms his day-to-day tasks. To demonstrate these methods the authors have trained an ANS network to drive a vehicle through simulated freeway traffic. (5 refs.)

GIRICO

Groupe Interuniversitaire de Recherche en Informatique Cognitive des Organisations

Sponsored by :

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

Workshop chairman

Gérard Simian Univ. Laval (C)

Local Arrangement Chairman

Yves Hudon GIRICO (C)

Program Committee Chairman

Bernard Moulin Univ. Laval (C)

Program Committee

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

ico '89

June 12 to 15, 1989 **Québec city, Canada**

INTERNATIONAL WORKSHOP ON COGNITIVE INFORMATICS APPLIED TO ORGANIZATIONS

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

Early Announcement

INVITED CONFERENCES

La représentation des procédures: le diagnostic, l'acquisition et l'utilisation des connaissances procédurales by Frederiksen (Can.)
LOQUI a natural language interface to databases by Binot (Belgium)
Analyse automatique de l'arabe by Mili (Tunisia)
L'intelligence artificielle dans les organisations, bilan et perspectives by Lorimy (Fr.)
La dynamique des application IA en entreprise: Perspectives by Haton (Fr.)
Artificial intelligence and Finance by Ross (USA)

COMMUNICATIONS

Approximately 30 communications will be presented in the following streams by researchers from Europe and North America:

- Knowledge representation and knowledge acquisition
- Object-oriented systems
- Planification and reasoning
- Intelligent interfaces
- Inferences and resolution
- Procedures, regulations and law
- Systems design aids
- Knowledge-based systems application

The workshop will include:
tutorials (June 12), industrial conferences (June 13), technical conferences (June 14 and 15).
An exhibition will be held concurrently with the workshop. Commercial products (hardware and software) and research prototypes will be exhibited.

For information : **Prof. Gérard Simian**
Colloque ICO'89
Université Laval dépt. d'informatique
Ste Foy, Québec G1K 7P4 Canada
phone: (418) 656-2312 or 656-7979

Upcoming Conferences

In Canada

First International Conference on Principles of Knowledge Representation and Reasoning
15 - 18 May 1989, Toronto, Ontario
See announcement on pages 58 - 59 of October 1988 issue for details.

6th Canadian Symposium on Instructional Technology
3 - 5 May 1989, Halifax, Nova Scotia
Focus: Computer-assisted learning (CAL) — theory and reality. Topics include: AI in education training; Innovations in instructional technology and courseware development; Technology transfer from researcher to user; Impact of expert systems and AI on CAL; User perspective on CAL. Tutorial sessions are on May 1-2.
Contact: F. Kewley, Sixth Canadian Symposium on Instructional Technology, Conference Services Office, NRC Canada, Ottawa, Ontario, K1A 0R6. Phone: (613) 993-9009. Telex: 053-3145.

Workshop in Game-tree Search
28 - 31 May 1989, Edmonton, Alberta
This workshop is part of the Canadian Information Processing Society's Annual Congress (CIPS). As well as presentation of papers, the 6th World Computer Chess Championship will be held, and hardware/software from all major manufacturers and research labs will compete in public competition. Topics include: Search algorithms; Search analysis; Parallelism; Planning; Learning; Knowledge representation and acquisition; Program design; Databases; Historical perspectives.
Submission material: 4 copies of full papers of approx. 4000 words for publication in proceedings; or 1500 word extended abstract for oral presentation. Submission deadline: Mar 1, 1989. Contact: Dr. T. Marsland, Computing Science Dept., U. of Alberta, Edmonton, AB T6G 2H1. Phone: (403) 432-3971.
Email: tony@alberta.cdn or ... .uucp.

2nd International Conference on AI and Law
13 - 16 June 1989, Vancouver, British Columbia
Topics include: Legal expert systems; Conceptual information retrieval; Case-based reasoning; Analogical reasoning; Representation of legal knowledge; Computational models of legal reasoning. Also invited are papers on relevant theoretical issues in AI (e.g., mixed paradigm systems using rules and cases) and in jurisprudence/legal philosophy (e.g., reasoning with precedents and rules).
Contact: Edwina Rissland, Dept. of Comp. and Info. Sci., U. of Massachusetts, Amherst, MA 01003. Phone: (413) 545-0332. Email: rissland@cs.umass.edu.

27th Annual Meeting of the Association for Computational Linguistics
26 - 29 June 1989, Vancouver, British Columbia
Topics include: Pragmatics, discourse, semantics, syntax, and the lexicon; Phonetics, phonology, and morphology; Interpreting and generating spoken and written language; Linguistic, mathematical, and psychological models of language; Machine translation and translation aids; Natural language interfaces; Message understanding systems.
Contact: Julia Hirschberg, ACL89 Program Chair, AT&T Bell Laboratories, 2D-450, 600 Mountain Ave., Murray Hill, NJ 07974. Phone: (201) 582-7496.
Email: julia@btl.att.com.

In the United States

Human Vision, Visual Processing, and Digital Display
19 - 20 January 1989, Los Angeles, California
This meeting is a subset of the SPSE/SPIE Symposium on Electronic Imaging, January 15-20, 1989.
Topics include: Models for human and machine vision; Color vision and color coding; Digitization, spatial sampling, and anti-aliasing; Vision-based algorithms for image processing; Psychophysics of image quality; Spatial/color/temporal interactions in perception and coding. Contact: Dr. Bernice Rogowitz, IBM T.J. Watson Research Center, PO Box 218, Yorktown Heights, NY 10598. Phone: (914) 945-1687. Email: rogowtz@ibm.com.

5th IEEE Conference on AI Applications
6 - 10 March 1989, Miami, Florida
Focus: Application of AI techniques to real-world problems. Topics include: Knowledge acquisition; Task-specific knowledge representation; Task-specific reasoning; Verification and validation; Diagnosis; Intelligent interfaces. Papers should focus on principles or case studies (in science, medicine, law, business, engineering, manufacturing, robotics).
Contact: Mark Fox, Robotics Institute, Carnegie Mellon U., Pittsburgh, PA 15213. Phone: (412) 268-3832.
Fax: (412) 268-5016. Telex: 854941.
Email: msf@isl1.ri.cmu.edu.

Symposium on Chinese Text Processing
16 - 17 March 1989, Boca Raton, Florida
Topics include: Character input and display; Character encoding; Design of terminals and keyboards; Chinese language lexicon; Chinese language parsing; Machine translation; Optical character recognition; Computer ideographics in the 90's; Speech recognition.
Contact: Dr. G. Kostopoulos, 1989 Symposium on Chinese Text Processing, Florida Atlantic U., Dept. of Comp. Engineering, 500 NW 20 St., Boca Raton, FL 33486. Phone: (407) 393-3463.

IEEE Workshop on Visual Motion

20 - 22 March 1989, Irvine, California

Focus: The representation and analysis of motion in image sequences. Topics include: Motion detection mechanisms; Optical flow and motion correspondence; Structure from motion; Event recognition and representation; Temporal planning and inferences; Control structures for dynamic scene analysis; Uncertainty in dynamic scene analysis; Applications in navigation, object manipulation and recognition.
Contact: Ellen Hildreth, AI Laboratory, 545 Technology Square, Cambridge, MA 02139.

AAAI Spring Symposium

28 - 30 March 1989, Stanford, California

Symposium titles include: AI and software engineering; AI in manufacturing; Knowledge system development tools and languages; Planning and search; Robot navigation; AI and limited rationality; Representation and compilation in high performance theorem proving; Spoken language systems; Innovative applications.
Contact: AAAI, 445 Burgess Dr., Menlo Park, CA 94025. Phone: (415) 328-3123.

Defeasible Reasoning with Specificity and Multiple Inheritance

7 - 9 April 1989, St. Louis, Missouri

This workshop is for those already active in the field. Reasoning systems include systems of defeasible, default and non-monotonic reasoning, alterations to conditional logic, non-standard logics, and systems of inheritance. The focal dispute is over what additional behavior a system should exhibit, given that it will prefer more specific defeasible rules when there is conflict.
Submission material: Short statement justifying attendance, indicating what issues/problems the individual intends to raise; and/or an extended abstract, 2 - 4 pages. Submission deadline for attendance justification and abstracts: February 14, 1989.
Contact: Ronal Loui, Dept. of Comp. Science, Washington U., St. Louis, MO 63130.

Modeling and Simulation

4 - 5 May 1989, Pittsburgh, Pennsylvania

Topics include: AI; Expert systems; Robotics and control theory; Microprocessors in computer science and education; Social, economic, geographic, regional science, and global modeling and simulation. Submission materials: 2 copies of 50 word abstract, plus summary. Submission deadline: January 31, 1989.
Contact: W. Vogt, Modeling and Simulation Conference, 348 Benedum Engineering Hall, U. Of Pittsburgh, Pittsburgh, PA 15261.

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

9 - 11 May 1989, Dallas, Texas

Topics include: AI applications and instruction; Intelligent tutoring systems; Knowledge acquisition and representation; Student modeling and cognitive diagnosis; Human computer instruction; Computational models of reasoning and learning; Evaluation of learning environments; Knowledge-based CAI systems; Authoring systems.
Contact: Prof. H. Maurer, IIG, Schiesstattgasse 4a, A-8010 Graz, Austria. Phone: 0043-316-70255/12.
Email: maurer@btx.uucp.

ICGA-89 3rd Conference on Genetic Algorithms

4 - 7 June 1989, Washington, District of Columbia

Topics include all aspects of genetic algorithms, including: Foundations of genetic algorithms; Search, optimization, machine learning using genetic algorithms; Classifier systems, Apportionment of credit algorithms; Relationships to other search and learning paradigms. Papers discussing specific applications (e.g., OR, engineering, science, etc.) are encouraged.
Submission material: 4 copies (hardcopy only) of full paper. Submission deadline: February 10, 1989.
Contact: Dr. J.D. Schaffer, Philips Laboratories, 345 Scarborough Rd, Briarcliff Manor, NY 10510. Phone: (914) 945-6168. Email: ds1@philabs.philips.com.

Computer Vision and Pattern Recognition

4 - 8 June 1989, San Diego, California

Topics include: Image processing; Pattern recognition; 3-D representation and recognition; Motion; Stereo; Visual navigation; Shape from (shading, contour, etc.); Vision systems and architectures; Applications of computer vision; AI in computer vision; Robust statistical methods in computer vision. Contact: Prof. R. Chellappa, Dept. of EE-Systems, U. of Southern California, Los Angeles, CA 90089-0272.

Industrial and Engineering Applications of AI and Expert Systems

6 - 9 June 1989, Tullahoma, Tennessee

Topics include: Knowledge acquisition and representation; Machine learning; Qualitative and temporal reasoning; Uncertainty management; Neural nets; Fault monitoring and diagnosis; Process control; Sunsor fusion; Human-machine interfaces; Vision; Pattern recognition; Natural language and speech processing; Robotics; Intelligent tutoring.
Contact: Dr. M. Ali, U. Of Tennessee Space Institute, Tullahoma, TN 37388. Phone: (615) 455-0631 ext. 236.

4th Rocky Mountain Conference on AI

8 - 9 June 1989, Denver, Colorado

Theme: Augmenting Human Intellect by Computer. Topics include: Intelligent support of human communication; Computer-supported cooperative work; Automated reasoning and problem solving; User interfaces and user interface management systems; Tutoring, training and education; Design, manufacturing and control; Planning; Human problem solving.
Contact: James Alexander, RMCAI Program Chair, U S WEST Advanced Technologies, 6200 S. Quebec #320, Englewood, CO 80111.

IASTED Conference on Expert Systems

12 - 14 June 1989, Los Angeles, California

For topics, see similar conference in Zurich, 26-28 June 1989. Submission deadline: Sept 15, 1989.
Contact: Expert Systems LA, IASTED Secretariat, PO Box 25, Stn. G, Calgary, AB, Canada T3A 2G1. Phone: (403) 270-3616. Fax: (403) 270-8855. Telex: 03-826670.

6th Workshop on Machine Learning

29 June - 1 July 1989, Ithaca, New York

The workshop will be divided into approximately 5 disjoint sessions, each focusing on a different theme and consisting of 50 participants invited on the basis of abstracts submitted to the session chair.
Contact: Alberto Segre, Dept. of Comp. Science, Cornell

U., Upson Hall, Ithaca, NY 14853-7501.
Phone: (607) 255-9196.
Email: ml89@cs.cornell.edu or segre@gvax.cs.cornell.edu.

Summer Computer Simulation Conference
24 - 27 July 1989, Austin, Texas
Topics include: Knowledge-based simulation theory; Intelligent simulation systems; Knowledge-based simulation tools; Knowledge-based systems (KBS) using simulation; Knowledge representation for simulation; Intelligent simulation control architectures; Applications of simulation techniques to KBS; Interactions between conventional simulations and KBS.
Contact: Society for Computer Simulation, ATTN: Group XIII, PO Box 17900, 4838 Ronson Court, Suite 'L', San Diego, CA 92117-7900.

CIE: Computers in Engineering Conference
30 July - 2 August 1989, Anaheim, California
Technical papers are invited in all areas relevant to the utilization of computers in the engineering profession; from research and development to applications, education, business and management issues and challenges.
Contact: (for AI, ES, KBS, Design Theory) Dr. G. Gabriele, Rensselaer Polytechnic Institute, Dept. of Mechanical Engineering, 110 Eight St., Troy, NY 12180-3590. Phone: (518) 276-2601.
(for Computers/Robotics in Education, Teaching CAD, Computer Aided Learning Systems) Dr. K. Tamma, Mechanical Engineering Dept., U. of Minnesota, 111 Church St., Minneapolis, MN 55455. Phone: (612) 625-1821.

4th Conference on Computers and Philosophy
10 - 12 August 1989, Pittsburgh, Pennsylvania
The conference will focus on intellectual, ethical and practical problems in all aspects of the relationships between philosophy and computers. Papers on AI are especially invited. There will be a showcase of instructional software for use in philosophy classes.
Submission deadline: March 31, 1989.
Contact: Leslie Burkholder, Center for the Design of Educational Computing, Carnegie Mellon U., Pittsburth, PA 15213. Email: lb0q@andrew.bitnet.

IJCAI 11th Joint Conference on AI
20 - 26 August 1989, Detroit, Michigan
The technical program consists of a "paper track" and a "videotape track". The paper track focuses on empirical, analytical, theoretical, conceptual, foundational aspects and applied research. The videotape track focuses on applications in all subfields best suited for this type of presentation.
Topics include: AI tools and technologies; Machine architectures, languages, shells; Search methods; Knowledge acquisition, learning, analogy; Real-time performance; Parallel and distributed processing; Cognitive modeling; Planning, scheduling, and reasoning about actions; Natural language, speech understanding and generation; Perception, vision, robotics; ITS; Design, manufacturing, control; Philosophical foundations, perspectives and attitudes, social implications.
Contact: IJCAI 89, c/o AAAI, 445 Burgess Drive, Menlo Park, CA 94025-3496. Call for workshop proposals are due Feb 1, 1989. Contact: J. Katz, MITRE Corp., MS L203, Burlington Rd., Bedford, MA 01730. Phone: (617) 271-5200. Fax: (617) 271-5161.
Email: katz@b.mitre.org.arpa

Outside North America

4th European Chapter Conference Assn. for Computational Linguistics
10 - 12 April 1989, Manchester, England
Topics include: Morphology; Knowledge representation and expert systems; Computer-assisted language learning; Machine translation; Lexical semantics; Computational models for the analysis and generation of language; Speech analysis and synthesis; Computational lexicography and lexicology; Syntax and semantics; Discourse analysis; Computational aids to translation; Natural language interfaces.
Contact: Harold Somers, Centre for Computational Linguistics, UMIST, PO Box 88, Manchester M60 1QD, England.

MIV-89: Workshop on Industrial Applications of Machine Intelligence and Vision
10-12 April 1989, Tokyo, Japan
Topics include: Knowledge representation and inference; Spatial reasoning; Vision algorithms; Learning and knowledge acquisition; New-generation programming; Expert systems; Human interface; Pattern information processing; Factory automation; Robotics; Multi-media database; Hardware architectures for AI and/or vision systems.
Contact: Prof. M. Ishizuka, Institute of Industrial Science, U. of Tokyo, 7-22-1, Roppongi, Minato-ku, Tokyo 106, Japan. Phone: +81-3-402-6231 ext. 2651.
Fax: +81-3-402-5078. Telex: 02423216 IISTYOJ.
Email: (Mr. H. Dohi) dohi%iis.u-tokyo.junet@relay.cs.net.

AISB '89
17 - 21 April 1989, Sussex, England
Held by the Society for the Study of Artificial Intelligence and Simulation of Behaviour (AISB).
Topics include: Knowledge acquisition and representation; Automated reasoning; Vision; Cognitive modeling; Commonsense reasoning; Learning; Psychological, philosophical or social implications; Search; Planning.
Contact: Dr T. Cohn, Dept. of Comp. Sci., U. of Warwick, Coventry CV4 7AL, UK. Email: agc@uk.ac.warwick.cs, or agc%uk.ac.warwick.cs@nss.cs.ucl.ac.uk.

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

International Symposium on Multiple-Valued Logic

29 - 31 May 1989, Guangzhou, China

Topics include: Automated reasoning; Algebraic and formal aspects; Circuit/device implementation; Fault detection and diagnosis; Logic design and switching theory; High speed computation; Probabilistic and variable-valued systems; Optical computing; Fuzzy logic; Philosophical aspects.
Contact: Prof. J. Muzio, Comp. Sci. Dept., U. of Victoria, PO Box 1700, Victoria, B.C., Canada V8W 2Y2.

Second Generation Expert Systems

29 May - 2 June 1989, Avignon, France

Second Generation ES are able to combine heuristic reasoning with deeper reasoning, based on a model of the problem domain. Topics include: Integration of different reasoning techniques; Architecture for combining heuristic reasoning and model-based reasoning; Cooperation of multiple expertise; Application of cooperative reasoning to real-world problems; Use of qualitative, causal or temporal reasoning techniques to augment heuristic reasoning; Integration of qualitative and quantitative reasoning.
Contact: Jean-Marc David, Laboratoires de Marcoussis, route de Nozay, 91460 Marcoussis, France.
Phone: 33-1-64.49.14.89. Fax: 33 - 1 - 64.49.06.94.

SCAI '89: 2nd Scandinavian Conference on AI

13 - 15 June 1989, Tampere, Finland

Topics include: Logic and AI theory; Knowledge representation and inference methods; Knowledge-based systems; Natural language and speech; AI tools and environments. An exhibition of AI tools and literature will be presented.
Contact: SCAI '89, Tampere U. of Technology, Ms. Raili Siekkinen, PO Box 527, SF-33101 Tampere, Finland.
Phone: Int +358 31 162441. Telex: 22313.
Email: scai89@tut.fi.

IASTED Conference on Expert Systems

26 - 28 June 1989, Zurich, Switzerland

IASTED is the International Association of Science and Technology for Development. Topics include: Knowledge acquisition and representation; ES design procedures and tools; Logic programming; Reasoning; Advisory systems; User interface; Testing and maintenance; Neural networks; and all areas of application. Special one day conferences on: ES in banking; ES in medicine, bioengineering and health care.
Submission material: Regular, short, survey and tutorial papers, 3 copies of 250 word abstract. Submission deadline: February 1, 1989.
Contact: Expert Systems, A. Kopp, IASTED, PO Box 354, CH-8053 Zurich, Switzerland. Telex: 817491.
(also see December conference in Los Angeles)

EKAW89: 3rd European Knowledge Acquisition for Knowledge-Based Systems Workshop

3 - 7 July 1989, Paris, France

Topics include: Elicitation/modeling of expertise (automatic, manual); Apprenticeship, explanation-based, other learning systems and integration of such systems with other knowledge acquisition techniques and systems (hypermedia, DBM systems, spreadsheets, simulators); Issues in cognition and expertise that affect the knowledge acquisition process; Extracting and modeling of knowledge from text; Knowledge acquisition methodology and training; Validation of knowledge acquisition techniques.
Submission material: 5 copies of extended abstract (8 pages) or full-length paper (20 pages max). Submission deadline: January 25, 1989.
Contact: B. Gaines, Dept. of Comp. Science, U. of Calgary, 2500 University Dr. NW, Calgary, AB, Canada T2N 1N4.
Phone: (403) 220-5901. Email: gaines@calgary.cdn.

4th Conference on AI in Engineering

10 - 13 July 1989, Cambridge, England

One major theme will be AI in design.
Contact: AIENG89, Computational Mechanics Institute, Ashurst Lodge, Ashurst Southampton SO4 2AA, England.
Email: (John Gero, Australia) john@archsci.su.oz.

IEEE ICIP '89 Image Processing Conference

5 - 8 September 1989, Singapore

Topics include: AI vision techniques; VLSI implementation; Image restoration/enhancement; Machine vision; Video communications; Office image processing; Image pattern recognition; Biomedical imaging; Remote sensing; System architecture.
Contact: Technical Program Chairman, ICIP '89, c/o Meeting Planners, 100 Beach Road, #33-01, Shaw Towers, Singapore 0718, Republic of Singapore.
Email: (Dr. Cho-Huak TEH) eletehch@nusvm.bitnet, or chteh@nuseev.bitnet.

CSCSI/SCEIO Application and Order Forms

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

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

CSCSI/SCEIO Membership
Canadian Artificial Intelligence Back Issues

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

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

Name .....

Mailing Address .....

.....

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

Computational Intelligence Subscriptions

( ) Please enter my subscription to *Computational Intelligence* at the CSCSI/SCEIO non-institutional member discount rate (Cdn$20).‡

Name .....

Mailing Address .....

.....

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

CSCSI/SCEIO Conference Proceedings

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

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

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

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

Name .....

Mailing Address .....

.....

National Research Council Canada
Conseil national de recherches Canada

Computational Intelligence

AN INTERNATIONAL JOURNAL

Computational Intelligence publishes, in English or French, high quality original theoretical or experimental research in computational (artificial) intelligence, including papers in areas of knowledge representation • natural language understanding • computational vision • applications of artificial intelligence • logic programming • theorem proving • learning • cognitive science • problem solving and planning • languages for artificial intelligence • speech understanding • game playing • philosophical implications • foundations of artificial intelligence •

Information for contributors:

Manuscripts for **Computational Intelligence** should be sent to the Editors, Computational Intelligence, School of Computing Science, Simon Fraser University, Burnaby, British Columbia V5A 1S6. High scientific standards are crucial and the journal encourages submissions of readable, issue-oriented presentations accessible to a general artificial intelligence audience.

Subscription Information:

Subscription information can be obtained (including special rates for members of the Canadian Society for Computational Studies of Intelligence) by calling or writing Stephen Prudhomme, General Manager, Research Journals, National Research Council of Canada, Ottawa, Ontario, Canada K1A 0R6 (613) 993-0362

Contents of Volume Three

During publication of Volume 3 **Computational Intelligence** introduced a new *Taking issue* section with an outstanding contribution debating logical approaches to artificial intelligence. This first "taking issue", edited by Hector Levesque, is a position paper by Drew McDermott entitled *A critique of pure reason.* The section includes commentaries by J. Allen and H. Kautz • D. Bobrow and M. Stefik • K. Bowen • R. Brachman • E. Charniak • J. deKleer • J. Doyle • K. Forbus • P. Hayes • C. Hewitt • R. Kowalski • R. Moore •G. Hinton • J. Hobbs • D. Israel • J. McCarthy and V. Lifschitz • N. Nilsson • S. Pentland •D. Poole • R. Reiter • S. Rosenchein • L. Schubert • B. Smith • M. Stickel and M. Tyson • R. Waldinger • T. Winograd • W. Woods, as well as a final response from McDermott.

Canada

AI System Development Tools

"Q&A"
Database/word processing package with an extensive Natural Language front end. Allows development of databases and documents using natural language. Can interface with popular 4th generation systems such as dBASE, WordStar, and Lotus 1-2-3. Over 50,000 copies sold in the U.S. since its introduction 2 years ago. $340.

NeuralWorks Professional II
Advanced, easy-to-use "graphic-intuitive" approach results in a seamless interface. User-defined neurodynamics, and loadable control strategies provide complete freedom in designing multi-paradigm networks or developing completely new network types. IBM PC AT or PS-2 and MAC Plus, MAC SE, MAC II $1,245. SUN/3 or SUN/4, $3,745.

Nexpert
Full scale expert system development tool for PCs with rich graphics support. Object hierarchy with multiple inheritance. Rules support pattern matching, integrated forward and backward chaining, automatic goal generation, non-monotonic reasoning. Open architecture allows integration with external programs. For use on IBM AT, 386, PS/2 and Macintosh II, $6,250. For use on DEC VAX 2000, II, III, under VMS/UIS, $10,000.

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

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

LPA PROLOG Professional
First PROLOG compiler to allow development of large applications on IBM PC AT. Provides a general purpose high-level programming language with a superb range of functions and data handling capabilities. PC Dual Compiler System $995. flex (Forward-chaining Logical EXpert) system toolkit (ms-dos) $495. MacProlog Color Edition $745.

ACCENT SA
A high quality text-to-speech convertor converts ordinary ASCII text into intelligible speech. Can be plugged directly into any IBM PC, PC/XT, PC AT or compatibles. Can speak in two modes: the text mode and the spell mode. Speech output may be directly connected to a speaker. A standard RS-232C serial port is provided for interface with a computer or terminal as the host. Prices range from $505—$1,480.

Arity Knowledge Systems Development Pack
A new generation AI prototyping tool for PCs. A prolog inerpreter/compiler, a rule/frame-based Expert Systems shell with object hierarchy and inheritance, the external file and database access, a screen design kit, linkage to conventional languages, and an IBM SQL access. $1,340.

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

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

Delphi Common Lisp (DCL)
Powerful new standard object-oriented programming facility. Compact, portable implementation of the complete Common Lisp standard. "Multithread" facility to create and control concurrent threads of program execution. Binary license for Sun 2, 3/50 and 3/100, $999. Sun 3/60, 3/200, and 386i, $1,995. Sun 4, $2,995. Apollo DN, $1,995.

DENDROS-1G Evaluation Board
All-analog, electronic ASIC and later VLSI implementation of a neural network which learns and categorizes in real-time. The dynamics of learning and sleep can be observed. $870.

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

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

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