Canadian Artificial Intelligence
Intelligence Artificielle au Canada

July/juillet 1990 No. 24

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

NeuralWorks Professional II and Designer Pack... the sensible approach to neural network development

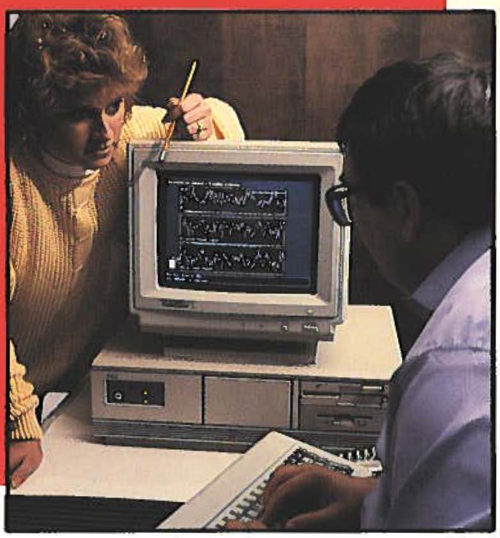

NEURALWARE, INC.

NeuralWorks Professional II, in conjunction with Designer Pack makes all other methods of neural development obsolete. The development environment (NeuralWorks Professional II) is seamlessly integrated with the deployment environment (NeuralWorks Designer Pack). It is easy to move back and forth between the two of them. You can try out as many networks and variations as you want. And when you are ready, you have "C" program modules with a well defined interface which will run identically to the network you prototyped.

NeuralWorks Professional II with Designer Pack ... the logical alternative.

Together with NeuralWorks Professional II, Designer Pack is the most elegant method for designing, testing, and deploying neural networks. The development blends naturally with application production. There is no need to learn special "Neural Network Programming Languages" which may be restricted to one or two special hardware systems. Build a network in a way that makes sense: graphically. Prototype your interfaces with User I/O. Turn it all into usable application code Designer Pack. Nothing wasted. No re-design, No re-learning, Simple, Elegant, Seamless...

The sensible approach to neural network development.

NeuralWorks Professional II and NeuralWorks Designer Pack are available on the Sun 3,4, 386i, N-Cube, MacIntosh, PC AT, XT, PS/2 and compatible machines.

User-Defined Neuro-Dynamics

Neurodynamics is the study of the generation and propagation of synchronized neural activity in biological systems.

User-Defined Neuro-Dynamics (UDND) is an advanced optional extension of NeuralWorks Professional II that allows users to write their own Summation, Transfer, Output, and Error functions, as well as Learning rules, in the "C" language.

These user-defined routines can be compiled and linked with NeuralWorks libraries to create a new customized NeuralWorks Professional II.

The UDND requires Zortech C version 1.07 or higher, and a Microsoft linker version 3.65.

APPLIED AI SYSTEMS, INC.

Applied AI Systems, Inc. (AAI) is an authorized dealer of NeuralWare, Inc in Canada and Asia-Pacific. In addition to sales and marketing o NeuralWare's products, AAI offers the following services:

Product demonstration: NeuralWorks can be demonstrated at AAI ir Ottawa, Canada. Also, a set of demo diskettes may be purchased from AAI for in-house demonstrations.

Introductory seminars: AAI conducts free introductory seminar on NeuralWorks and other AI products from time to time. Call AA for information on the next seminar, or to register your name on AAI' mailing list.

Training: AAI conducts hands-on training in NeuralWorks and other A products at its training facility in Ottawa, or at a client's site. Contact AA for the details.

After-sales support: A support package is available for NeuralWorks The package includes upgrades to the software and technical support b telephone. Registered customers may also visit AAI for technica consultation.

Consulting: As a full service AI company, AAI's staff can provid consultation services to governmental and industrial customers. Th company has been involved in applicable research of various subfield of AI since 1983.

Application Development: Applicability of neural network technolog is already widespread. The following list covers some of the areas wher an application has already been fielded or is in development:

• robotics (perception, control, action fusion, generation of actio sequence) • artificial skin/muscle, vision and scene analysis • characte and object recognition • process control (process monitoring, qua ity control, generation of flexible manufacturing schedule) • exper systems (diagnosis, control expert system) and automated know edge acquisition • voice recognition • signal processing (intelligen noise filtering, signal formatting) • radar clutter elimination • ris analysis for credit authorization • prediction • detection o explosives • autonomous vehicles (sensory and control processes map generation, map interpretation, navigation an maneuvering) • various multiple constraint satisfaction problem (urban traffic optimization, communication network optimization amenity optimization) • recognition (sonar images, on-orbit sate lites, mineral ore discovery, radar return signals, radar signal sourc identification, signature authentication).

AAI will assist business/industrial clients in the development of application using neural networks.

For PC from $1,825, Sun 3 and 4 from $2,440

NeuralWare
INCORPORATED

Phone or write for more information

Applied AI Systems, Inc.
Gateway Business Park
340 March Road, Suite 500
Kanata (Ottawa), Ontario, Canada K2K 2E4
Tel. (613) 592-3030 • Fax. (613) 592-2333

Canadian Artificial Intelligence

Intelligence Artificielle au Canada

No. 24 July/juillet 1990

Contents

Canadian Society for Computational Studies of Intelligence

Founded in 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 en 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 dernière page. 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

President/Président: Ian Witten, Head of Computer Science, U. of Calgary, Calgary, AB T2N 1N4, 403-220-6780; email: ian@cpsc.UCalgary.ca
Past-President / Président Précédent: Dick Peacocke, Bell-Northern Research, Box 3511, Station C, Ottawa, ON K1Y 4H7, 613-765-2629; BITNET: richard @bnr.ca
Vice-President/Vice-Président: Janice Glasgow, Dept. of Computing and Information Science, Queen's U., Kingston , ON K7L 3N6, glasgow@qucis.QueensU.ca
Secretary/Secrétaire: Peter Patel-Schneider, AT&T Bell Laboratories, 600 Mountain Ave., Murray Hill, NJ 07974, USA, 1-201-582-3399; email: pfps@research.att.com
Treasurer/Trésorier: Grant Thomas, Partner, Price Waterhouse, Suite 1100, 180 Elgin St., Ottawa, ON K2P 2K3, 613-238-8200
Managing Editor/Editeur: Roy Masrani, Alberta Research Council, 6815, 8th St. N.E., 3rd Floor, Calgary, AB T2E 7H7, 403-297-2656; masrani@noah.arc.ab.ca

Canadian Artificial Intelligence

Intelligence Artificielle au Canada

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

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

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

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 soumissions ne doivent 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 Écrivez à l'addresse à la gauche ou téléphonez 403-297-2600.

Please send submissions to / Prière d'envoyer contributions à:
masrani@noah.arc.ab.ca
or to / ou à: Roy Masrani
Canadian Artificial Intelligence
Alberta Research Council
6815 8th Street NE, 3rd floor
Calgary, Alberta, CANADA T2E 7H7

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

Magaging Editor/Editeur: Roy Masrani
Editor Emeritus/Rédacteur emeritus: Graeme Hirst
Translation/Traduction: Anne Parent, Diane Lefebvre, Raymond Aubin
Production Manager: Carol Tubman
Production: Ad Ventures Studio Inc.
Advertising/Réclame: Marco Ariano
Publications/Publications: Graeme Hirst

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

COMMUNICATIONS
COMMUNICATIONS

Parting Notes

First of all I'd like to thank the outgoing executive, Renato, Jan, Bill, and Gord, for all their work during our term in office. There have always been a hundred things needing to be done, and loose ends calling for attention. The pile of e-mail has been an ever-growing one!

Most of the work as CSCSI President has been satisfying and useful, I hope. The magazine has been in capable hands and continued to thrive, and we can all be proud of Computational Intelligence. We have had two excellent conferences (in Edmonton and Ottawa) during my four years of office, and we can look forward to CSCSI-92, as well as IJCAI-95 in Montreal.

Working with CIPS has had its ups and downs. Some of these administrative tangles seem to be being resolved, but we shall have to see. The new executive is a very able set of people, and I'm confident that they will give good direction to CSCSI, and handle this administrative issue effectively. I look forward to helping them do that, and I wish them luck in their new positions.

Dick Peacock, Past-President

New Executive

President — Ian Witten

Ian Witten received degrees in Mathematics from Cambridge University, Computer Science from the University of Calgary, and Electrical Engineering from Essex University, England. A Lecturer and subsequently Senior Lecturer at Essex University from 1970, he returned to Calgary in 1980 where he served as Head of Computer Science from 1982 to 1985.

The underlying theme of his current research is the exploitation of information about a user's past behavior to expedite interaction in the future. In pursuit of this theme he has been drawn into machine learning, which seeks ways to summarize, restructure, and generalize past experience; adaptive text compression, in other words using information about past text to encode upcoming characters; and user modeling, which is the general area of characterizing user behavior.

Ian has published widely on machine learning, text compression, autonomous systems, speech synthesis and signal processing, hypertext, and computer typography. He has produced four books: Communicating with Microcomputers (Academic Press, 1980), Principles of Computer Speech (Academic Press, 1982), Talking with Computers (Prentice Hall, 1986), and Text Compression (Prentice Hall, 1990), the last one co-authored with T. Bell and J. Cleary; and is currently working on another, the Reactive Keyboard, (to be published by Cambridge University Press), co-authored with J. Darragh.

Vice-President — Janice Glasgow

Janice Glasgow is an associate professor in the Department of Computing and Information Science at Queen's University. She received her B.Sc from the University of Alberta and her M.Math and Ph.D degrees from the University of Waterloo. Dr. Glasgow's current research interests include computational models for imagery, case-based reasoning and machine learning. She is a Principal Investigator in the Institute for Robotics and Intelligence Systems (IRIS) Federal Network of Centers of Excellence of Canada and a Co-investigator in the Information Technology Research Corporation (ITRC) and the Manufacturing Research Corporation (MRCO) Provincial Centers of Excellence of Ontario.

Secretary — Peter F. Patel-Schneider

Peter F. Patel-Schneider is currently employed by *AT&T* Bell Laboratories in the AI Principles Research Department. He received a Ph.D. from the University of Toronto in 1987 in the area of knowledge representation. He worked for the Fairchild Laboratory for Artificial Intelligence Research and its successor, Schlumberger Palo Alto Research, from 1983 to 1988. Peter's current research interests include non-standard logics for knowledge representation and designing, analyzing, and building terminological knowledge representation systems.

Treasurer — Grant Thomas

Grant Thomas is an associate and former partner of Price Waterhouse Canada. As a consultant in advanced technology, he has been heavily involved in the investigation of advanced information technologies and their potential roles in environment impact assessment and environmental management. Mr. Thomas chairs the National Research Council Associate Committee for Artificial Intelligence. He is also a member of the NSERC Strategic Grant Review Panel. He sits on the board of several information technology professional associations.

Dick Peacocke and Roy Masrani will both remain on the executive as Past President and Managing Editor of Canadian Artificial Intelligence Magazine respectively.

Letters to the Editor:

Dear Editor,

Last March I submitted a book review to Canadian AI's book editor. The submission was done electronically, and the submitted text had all french accents included in a way that would be non-ambiguous to anyone with even the most rudimentary knowledge of French. I had also included a printed version of the article, should the proofreader need to verify anything. Much to my dismay, however, the review as printed came out with totally mangled accents (with all instances of "a" with a grave accent mutated into "e" with a circumflex accent, and all instances of "e" with a grave accent appeared with acute accents instead).

Furthermore, there were more glaring typos in the title of the book in the review, and the short bibliographical notes. Another book review by Guy Lapalme had much the same fate.

I find it very hard to be indulgent; we aren't talking about an occasional slip of the fingers when transcribing text in an unfamiliar language — the submission was electronic, and a reference printout was available for proofreading. Many francophones may well interpret the lack of care shown in the inclusion of French text as a measure of the respect they are getting, and that is truly unfortunate.

Jean-François Lamy, Groupe Sobeco, Inc. Montre'al, Qc lamy@sobeco.com

[Thank you for your frank feedback. These are indeed embarrassing errors for which these is no excuse. Please accept my apologies; I will be especially diligent in the future. - RM]

AI NEWS
NOUVELLES DE L'IA

Trains run on time, thanks to AI

Anyone familiar with operations at New York's Grand Central Station — and especially those who insist on more accurately calling it Grand Central Terminal — will appreciate the scheduling problems encountered at one of France's busiest terminals, the Gare de l'Est in Paris. All train terminals share the same logistics problem: they have to funnel trains from many parallel platforms into a very few mainline tracks. At Gare de l'Est, 30 platform tracks funnel into just six mainline tracks, and some 1,100 trains a day must be hurried through this bottleneck. That's one train every 30 seconds during busy periods every day.

Raymond Moulin, manager of expert system projects at the Cybernetic and Technical Research Department of SNCF, France's railway authority, explained some of the problems that confront human schedulers. Traffic levels are near the theoretical maximum. Each train must be assigned one of 640 possible routes into and out of the station. Local and long distance trains are mixed on platform assignments. A single delayed train can cause a chain reaction during rush hours that reverberates through the schedule of subsequent trains as long as four hours afterward. And when a track or platform must be taken out of service for repairs, as many as 250 trains may have to be diverted each day.

Human expertise

Moulin pays tribute to the dispatchers who, like chess masters, have solved these knotty problems quickly every day for scores of years. "Only specialists at the Gare de l'Est have the skills to reroute these trains without creating delays — skills that derive from 10 to 15 years' experience of working at the station."

Now, however, those skills have been built into GESPI, an expert system running on a TI Explorer computer. The human experts now need intervene only to relax constraints beyond the system's authority to relax them. GESPI was tested at the Gare de l'Est for half a year in 1988, and has been in full operational use every day since early 1989. The system was developed by SNCF in collaboration with GSI-TECSI, a specialized AI company, using a TI Explorer computer and Carnegie Group's KnowledgeCraft development tool.

Moulin reports that GESPI is working well, and that clear-cut benefits have been achieved. "It has helped to improve the quality of plans for the daily movements of trains, and to prepare future timetables," he said. "The experience gained in this first project will be put to good use in similar implementations at other stations," he said. "GESPI has two goals: increasing the regular traffic flow, and facilitating the integration of new daily traffic conditions — supplementary trains and track repairs. There is no algorithmic solution to this problem."

The number of possible solutions to such problems is astronomical, as the result of "combinatorial explosion." As in chess, the combinations of all possible moves and counter-moves are so numerous that a powerful computer may take years to explore them all and select the best. Yet in the real world of railroading, a solution is needed in minutes.

Until now, the solution has been for the human experts to use rules-of-thumb — heuristics — to give them a working solution that may not be the absolute best, but that's adequate. These rules may be enunciated as constraints that say what may and what may not be done in terms of train routes, and the effects of any changed route on all the others. The basic rule in railroading, of course, is that two trains must not occupy the same space at the same time; the first corollary is that, on a single track, no train may pass another from either direction. These are rules that must never be violated; other rules range across a spectrum whose other extreme is "slightly desirable if convenient." At this extreme are such rules as "Don't assign a train to a platform until the previous debarkees have fully cleared the platform," which may be relaxed to solve problems.

The first advantage of GESPI is that a computer-based system can run through many more possible arrangements than humans can in a limited time, resulting in improved management of the terminal. But GESPI goes much farther. When extra trains are added, as for holidays, it often turns out that no solution is possible that obeys all the constraints. GESPI has metaknowledge that lets it begin to relax the least-crucial rules until satisfactory solutions are possible. Only when it has exhausted its authority to relax rules does GESPI call for human intervention to relax more-important rules to meet emergencies.

Relations, objects

GESPI employs an object-oriented approach. A plan's "elements" — trains, tracks and platforms — are linked via "relations." The system begins by examining a possible route, and as soon as a potential conflict between two trains occurs, the type of conflict triggers an appropriate rule that recommends how to remove it. Each time such an event happens, a relation is established between the conflicting objects, representing the type of conflict.

These processes quickly build a network of objects and their relation. But it isn't feasible to produce a complete solution all at once, so GESPI sub-divides the problem into small elements, and starts working on areas where experience has shown the trains to be most exposed to potential conflicts. "The rules make it possible to identify these situations in a practical way, without calculation, and very reliably. It is a proven method, because the rules have been used regularly for years," Moulin said.

GESPI normally starts in batch mode and applies its rule base without an operator present, listing any situations that cannot be resolved. Then, in the interactive mode, the human operator can intervene by using a graphical interface, relaxing constraints more than the system is permitted to do, until a final plan is produced. The plan reproduces the paper documents formerly used by the humans to produce their plans, so the need for behavioral changes outside dispatching was held to a minimum.

GESPI's interface also gives users a clear window into the system, so they can easily see the effects of manually modifying data in the system. And it invites them to investigate new possibilities off line. "The experts can modify and manipulate the documents they're used to on the Explorer's screen, but they can also exploit new methods of working. Whichever they do, they can use the system without knowing anything about the underlying AI techniques," Moulin said.

Even casual observers of the French railroading scene must be impressed with their system, which is one of the world's most advanced. While the fastest train in the US is the 125-mph Metroliner connecting Washington and New York, and the Japanese Shinkansen bullet trains travel at 150 mph, the French Trains a Grande Vitesse (TGV) Atlantique cruise at 186 mph. And the French have earlier TGV models in abundance. One would therefore have expected the French railway system to adopt expert systems technology early on: it's a great way to run a railroad.

Smart trains

Meanwhile, at Canadian National (Montreal, Quebec), the smart money has smart trains taking over the North American railway industry during the next decade. Smart trains are the prodigy of Advanced Train Control Systems (ATCS), the technology that many are betting will change railroading as profoundly as the switch from steam to diesel power almost a half century ago. Smart trains are said to be safer, more efficient, will improve productivity, will eliminate paperwork, and will diagnose their own mechanical problems.

After six years on the drawing board, ATCS is moving onto the rails — albeit cautiously — at several major carriers in Canada and the US. Canadian National (CN) recently turned on the first ATCS operating communications system in North America at one of two test-beds over 100 miles of the railway's main and secondary lines near Toronto. The system, covering CN's Halton, Oakville and Dundas subdivisions, is part of an industry-wide effort to demonstrate and prove ATCS concepts and system reliability.

At the same time, the process of identifying, rating and exploiting potential uses for the technology beyond its primary role of electronic rail-traffic manager is moving into high gear. As far as the planners are concerned, ATCS is just about on schedule, even though it took five years for the industry to agree on the final design and specifications.

"There isn't a whole lot of hardware on the ground yet," said John Reoch, CN Rail's assistant vice president, operations, "but in the context of the kind of technological breakthroughs that are involved, and the size and complexity of the project, not to mention the potential cost, I'd say it's just about on track."

Most railways will be watching pilot projects like the British Columbia North Line installation — a leading edge application which will break new ground for the railway industry, according to Walter Friesen, spokesman for a consortium of suppliers headed by the SEL Division of Alcatel Canada Inc. (Don Mills, Ontario). The consortium is running a prototype ATCS installation on CN's Hagersville subdivision south of Brantford to put hardware and integrated systems destined for the B.C. North Line through their paces. By 1991, the railway will have a first operational ATCS installation on 187 miles of the B.C. North Line between Prince George, B.C., and Jasper, Alberta.

Advanced Train Control is just as exciting for equipment suppliers because ATCS equipment is built to industry-wide specifications. What they can sell to one railway, they can sell to them all, or so they hope. For SEL and consortium partners Motorola Canada (North York, Ontario) and Vapor Canada (St. Laurent, Quebec), that represents a crack at a North American market worth billions of dollars.

To create a smart train, you put a computer on a locomotive, and have it "talk" via a data radio hookup to a central ATCS computer, which directs traffic for maximum safety and efficiency. The ATCS computer sends train movement orders to the onboard computer, which displays them for the crew.

The glue for ATCS is the data radio link. To determine the optimum location for radio base stations, CN Rail purchased a National Research Council software package that computes the signal strength needed to provide proper coverage from any given point along the B.C. North Line.

The B.C. system is designed to perform such ATCS functions as the receipt and display of train movement instructions in the locomotive cab by means of an on-board computer. From the safety viewpoint, enforcement of such instructions is the most exciting aspect of ATCS.

In a fully configured ATCS environment, the system can detect and override human error, virtually ruling out the possibility of train collisions. The ATCS computer will monitor, separate and pace traffic just like an air traffic controller, sending train movement orders to the on-board computer. If something goes awry — if the train is speeding or about to pass a mandatory stop — the computer can enforce speed limits or stop the train.

What's exciting to the railways is that the ATCS framework — a "live" data network linking central and remote computers as well as any trackside equipment — can serve a myriad of other masters. Once a train becomes "smart," those smarts can be put to work for the marketing department, for the motive power repair shops, for the people who allocate motive power and freight cars, and for the track repair gangs.

The on-board computer can have a diagnostic capability to monitor locomotive health and forward trouble reports to the next maintenance point. With improved train control, more trains can operate on a line. That may encourage railways to operate shorter, faster trains. And trains can be paced to maximize fuel efficiency. Stop-and-go driving burns more fuel.

Perhaps the leading business application so far is a computerized work order reporting system. Train consist information and work orders can be transmitted between the locomotive and a railway's business computers. That has the potential to greatly cut-down on paperwork and enable the railway to respond more quickly to customer pick-up and delivery requests.

CN and a joint venture company owned by the Union Pacific Railroad, Tandem Computers, and Alcatel SEL are trying out a work order reporting system on a 125 mile territory near Belleville, in eastern Ontario.

"Right now, the critical issue for ATCS is developing the business case," said John Reoch. "There has to be a whole range of business payback — return on investment — that complements the safety aspects of the system for the railway industry to make the kind of investments that are involved."

Over the next two to five years, "a lot of people will start installing ATCS, but there will also be some applications of conventional technology," said Walter Friesen. "In the seven to 10-year time frame, however, ATCS will be the technology of choice for North American railways." The current industry projection is for the installation of ATCS along 200,000 miles of track by the end of the century.

Reprinted with permission – AI Week, March 1, 1990.

Carnegie Group

(Pittsburgh, PA) has signed a contract for its Service Maintenance Planner (SMP) with the Canadian National Railroad. Canadian National, based in Montreal, will use the system to resolve conflicts between a railroad's train service and track maintenance schedules. The rail company operates essentially as a freight carrier throughout Canada on over 22,000 miles of track, although it works cooperatively with the government to carry passengers. SMP will ultimately be integrated into the railroad's Track Maintenance Program, a computerized system for managing all of Canadian National's track maintenance. Carnegie originally developed SMP for Burlington Northern Railroad, and claims that it has earned that company back its initial development price, with an expected annual savings of $25 million produced by the system by 1991. The cost for SMP is now $295,000, and includes a year of software support for the various KnowledgeCraft modules offered as part of SMP. The package runs on both the Sun-4 and DEC MicroVAX.

Carnegie has also enhanced its Operations Planner application for the PC/AT. The $3900 package, which is used for manufacturing planning and budgeting, now allows for a variety of cost scenarios as applied to specific work centers or manufactured products. The company also announced the availability of Knowledge Craft 3.3 tools such as CRL, CRL-OPS. and Allegro Common LISP for the DECstation — DEC's ULTRIX-based RISC machine. Priced from $4000 to $5000, the products help round out the DECstation as a low-cost delivery vehicle for a variety of AI applications.

Reprinted with permission — AI Trends, February, 1990.

STEAR Awards Five Contracts

The Space Station Program's Strategic Technologies for Automation and Robotics (STEAR) branch awarded the last of five Phase Two contracts September 21 for "Automation of Operations" on the Space Station's Mobile Servicing System (MSS).

The STEAR Program helps Canadian industry develop advanced technology, particularly automation and robotics, for use in the design of MSS. The new technology is expected to create terrestrial spinoffs and longterm industrial benefits.

The five teams who won the $900,000 contracts are:

- Artificial Intelligence Inc.; Atomic Energy of Canada Ltd., (B.C., Ontario)
- Softwords (Press Porcepic); Alberta Research Council (B.C., Alberta)
- Det Norske Veritas; Alberta Research Council (Alberta)
- Spectrum Engineering; Queens University (Ontario)
- Dynacon; N.B. Research and Productivity Council; CRS Plus; University of Toronto; University of Ottawa (Ontario, N.B.)

Artificial Intelligence Inc. and Softwords will develop plans involving the use of expert systems for training astronauts and technical personnel. Det Norske Veritas and Spectrum Engineering are investigating expert systems for fault diagnosis. The expert system uses a computer program which stores data about a machine's normal performance to diagnose faults in the machine by searching its memory. Dynacon is developing a sophisticated planner for operational applications.

The five contracts will be completed by 1992. Experts from the National Research Council of Canada and the Space Mechanics Group of the Canadian Space Agency will supply scientific advice and direction. "Automation of Operations" is being managed by Dr. Donald Smith of the STEAR Program.

MSS is a multi-armed robotic unit which will be used to assemble and maintain the Space Station, beginning in 1995-96.

STEAR Décerne Cinq Contrats

Une ramification du programme "Station spatiale" appellée (STEAR): Strategic Technologies for Automation and Robotics, à accorder le 21 septembre ses cinq derniers contrats (phase two) pour des projets se rapportant à "l'Automation des Opérations" sur le MSS (Mobile Servicing System) de la station spatiale.

Le programme STEAR aide les industries canadiennes à développer de la technologie d'avant-garde, en particulier dans le domaine de l'automation et de la robotique, qui sera utilisée dans la conception du MSS. On s'attend à ce que cette nouvelle technologie ai des applications secondaires sur terre et quelle soit, à long terme, bénéfique aux industries.

Voici les cinq équipes qui se sont vues décerner les contrats de $900,000 chacun:

- Artificial Intelligence Inc.; Energie Atomique du Canada Ltd. (B.C. et Ontario)
- Softwords (Press Porcepic), Conseil de Recherche de l'Alberta (B.C. et Alberta)
- Det Norske Veritas, Conseil de Recherche de l'Alberta (Alberta)
- Spectrum Engineering; Université Queens (Ontario)
- Dynacon; Conseil de Recherche et Productivité N.B.; CRS Plus; Université de Toronto; Université de Ottawa (Ontario et N.B.)

Artificial Intelligence Inc. et Softwords vont développer des projets dans les quels des systèmes experts seront utiliser pour la formation des astronautes et techniciens. Det Norske Veritas et Spectrum Engineering vont explorer la possibilité d'utiliser les systèmes experts dans le diagnostic des défaillances. Le système expert utilise un programme d'ordinateur qui met en réserve les données correspondantes au niveau de performance normal de la machine. Ceci permet de diagnostiquer les défaillances dans le fonctionnement de la machine en fouillant sa mémoire. Dynacon pour sa part developpera un planificateur très sophistiqué orienté vers des applications opérationnelles. Les cinq contrats seront achevés en 1992. Des experts, membres du Conseil de Recherche National du Canada et de l'Agence Spatiale Canadienne (Groupe des Mécaniciens de l'Espace), seront en charge de donner des conseils scientifiques et des directives. Le Projet "Automation des Opérations" est dirigé par Dr. Donald Smith du programme STEAR.

Le MSS est une unité robotique controlant plusieurs bras, qui sera utilisée pour assembler et entretenir la station spatiale. Le MSS sera en operation en 1995-96.

NEWS RELEASE

ISTC MINISTER ANNOUNCES FEDERAL FUNDING FOR ARTIFICIAL INTELLIGENCE RESEARCH AND DEVELOPMENT FUND

OTTAWA, April 5, 1990 — Industry, Science and Technology Minister Benoît Bouchard, announced today the funding of eight projects, totalling $2 million, under the auspices of the Artificial Intelligence Research and Development Fund. These projects are being undertaken by the Department of Communications, Energy, Mines and Resources Canada, Environment Canada, Health and Welfare Canada, Transport Canada, and the Department of National Defence, on a cost-shared basis. The majority of work on each project will be contracted out to firms in the private sector with capabilities in artificial intelligence.

The Artificial Intelligence Research and Development Fund is a procurement-based program which will use the federal government as a test bed to assist in the development of private sector capabilities in artificial intelligence. "Through the Artificial Intelligence Research and Development Fund, the federal government is taking the lead in using high risk advanced technologies, which also have the potential to improve the effectiveness of government operations. The Fund is a demonstration of ISTC's role in building a competitive economy in close partnership with innovative Canadian firms," Minister Bouchard said.

The Department of Supply and Services manages the process of contracting out the work on these systems according to its standard procedures. Each sponsoring department is responsible for the management of its own project. The Fund is a component of ISTC's Strategic Technologies Program, which assists Canadian industry in responding to the challenge of rapid technological change in information technology, biotechnology and advanced industrial materials.

"Using the federal government as a proving ground, Canadian firms will be able to explore the potential of artificial intelligence and demonstrate the considerable pay-offs which are expected," the Minister added.

The Artificial Intelligence Research and Development Fund will match the contribution of government departments and is a $10 million, five year program. This initial allocation from the fund, of approximately $2 million over four years, will result in a total of over $4.6 million being devoted to AI technology development. It is expected that, over the life of the program, total project investments for these new software technologies will be over $20 million because of the leverage exerted by the Fund.

LE MINISTRE D'ISTC ANNONCE L'OCTRI D'UNE CONTRIBUTION FÉDÉRALE AU FONDS DE RECHERCHE DÉVELOPPEMENT EN INTELLIGENCE ARTIFICIELLE

OTTAWA, le 5 avril 1990 - Le ministre de l'Industrie, des Sciences et de la Technologie, M. Benoît Bouchard, a annoncé aujourd'hui le financement de huit initiatives, s'élevant à 2 millions de dollars, sous les auspices du Fonds de recherche développement en intelligence artificielle. Ces initiatives seront mises en oeuvre selon la formule des frais partagés par les ministères suivants : celui des Communications, de l'Énergie, des Mines et Ressources, de l'Environnement, de la Santé et du Bien-être social, de même que celui des Transports. La majeure partie du travail, effectuée dans le cadre de chaque initiative, sera confiée en sous-traitance au secteur privé avec les possibilités de l'intelligence artificielle.

Le Fonds de recherche-développement en intelligence artificielle est un programme de marchés qui utilisera le gouvernement fédéral comme banc d'essai pour contributer à la mise sur pied de capacités en intelligence arificielle dans le secteur privé. « Par l'intermédiaire du Fonds de recherche-developpement en intelligence artificielle, le gouvernement fédéral prend l'initative d'utiliser des technologies de pointe à risque élevé, qui ont le potentiel d'améliorer l'efficacité des opérations du gouvernement. Le Fonds témoigne du rôle que ISTC joue pour bâtir une économic concurrentielle en étroite association avec des entreprises canadiennes innovatrices, » a déclaré le ministre Bouchard.

Le Ministère d'Approvisionnoment et Services gère la sous-traitance de ces systèmes selon les procédures normales. Tous les ministères qui financent sont tenus responsables de leur propre projet. Le Fonds est un volet du Programme des technologies stratégiques d'ISTC qui aide l'industrie canadienne à relever le défi posé par l'accélération de l'évolution technologique en technologie de l'information, en biotechnologie et en matériaux industriels avancés.

« En utilisant le gouvernement fédéral comme banc d'essai, les entreprises canadiennes seront en mesure d'explorer les possibilités de l'intelligence artifcielle et de montrer que l'on peut s'attendre à des bénéfices considérables, » a ajouté le Ministre.

Le Fonds de recherche-développement en intelligence artificielle est un programme de 10 millions de dollars, d'une durée de cinq ans. L'affectation initiale du Fonds, d'environ 2 millions de dollars étalé sur quatre ans, aura pour résultat d'attribuer plus de 4,6 millions de dollars à l'élaboration d'une technologie en intelligence artificielle. Au cours de la durée du Fonds, on s'attend à ce que le total des investissements destinés à ces nouvelles technologies, dans le domaine des logiciels, dépasse 20 millions de dollars.

How to get a research grant

Ian H. Witten

Aprés avoir été membre pendant trois années, du commité de selection des bourses pour ordinateurs et sciences informatiques du conseil de recherche du Canada pour les sciences naturelles et l'ingénierie (NSERC), l'auteur a enfin pu observer et apprecier le travail de recherche fait dans les diverses universités canadiennes. On realise que plusiers chercheurs talentieux ne reçoivent pas les fonds nécessaires à leurs recherches parce qu'ils ne savent pas bien rédiger leurs demandes de bourse. Votre demande doit être bien détaillée afin de permettre au commité de selection de décider de façon équitable.

Cet article s'adresse en particulier aux demandes de bourses pour les sciences informatiques du NSERC, mais s'applique aussi à d'autres bourses de recherche.

1. Introduction

I've just spent three years on the Natural Sciences and Engineering Research Council of Canada (NSERC) grant selection committee for Computer and Information Science. This is an arduous job, but a worthwhile — and very interesting — one. It provides an opportunity to see most of the computer science research going on in Canadian universities, and although you suffer from terrible information overload you do gain an appreciation for the breadth and excellence of the work being done. The most painful part of the job is the extraordinarily inadequate amount of money that granting agencies have to work with, and the need to reduce, or even cancel, funding for many worthwhile projects because of the extremely competitive nature of the awarding process and the dire shortage of funds.

The second most painful part of the job, which prompted me to write this article, is seeing how many capable researchers remain unfunded because they are unaware of how to write good research proposals. Many interesting projects go by the board because they are inadequately described. In the hotly competitive environment in which the grant selection committee operates, it is inevitable that inadequate or poorly-prepared research proposals receive little benefit of the doubt. The onus lies squarely on the applicant to provide clear evidence on which the committee can base a decision. This note summarizes what I have learned about how to write research proposals, through having had to evaluate a lot of examples — good and bad — over the past three years. Provided certain mistakes are avoided, the excellence of a proposal hinges on the originality and impact of the research, and this article won't help you with that! But there are some simple guidelines that must be followed to generate a well-presented proposal.

Three principal factors are taken into account when evaluating a research grant application:

- the quality of the research program;
- the quality of the written proposal;
- the quality of the researcher.

Ian Witten is a professor in Computer Science at the University of Calgary. He is also the new president of the CSCSI.

The first factor, the quality of research, is discussed in Section 2 below. The second, the quality of the proposal, is addressed head-on in Section 3. A researcher's reputation, which is built—up or down—over time, strongly influences how his or her proposals are seen, and Section 4 gives some advice on how to present yourself in the best light. Section 5 sketches how a grant selection committee actually works. Section 6 gives some information about refereeing research grant applications, an activity that—though often seen as a chore—is absolutely essential for the health of the discipline.

This article is targetted at proposals for NSERC computer science operating grants, which are intended to provide basic support for individual researchers' work—although the same general ideas apply to any research proposal. NSERC stresses longer-term funding for individual researchers' programs more than funding for particular projects. Other granting programs have different priorities, and this should be borne in mind when preparing proposals. It should be emphasized that the views expressed here do not necessarily reflect the official policy of NSERC or any other body. Bundy (1988) has written a useful note from which some of these ideas are derived.

2. Research ideas

To do research you must formulate a *question* that your work will strive to answer. This should not just be an isolated question, but one relating to a longer-term research *theme* that evolves over a substantial part of your career—certainly much longer than the 3-year term of the average research grant. Moreover you should begin with not just a single question, but a few (although not too many) that differ in riskiness, and hence potential value. You must be able to evaluate these research questions yourself, so that you can pick good ones and present them clearly.

2.1. Generating research questions

In computer science it should not be hard to come up with good research questions. The field is young and there is much to do. Technology changes constantly, radically altering the boundaries of what is feasible, and new possibilities for research are continually opening up. There are fertile opportunities in replicating previous work more systematically and in greater depth—rational reconstruction of programs, experimental evaluation and comparison, tightening up existing conceptual frameworks, and so on. There are plenty of avenues for research in computer science!

Nevertheless, it may still be difficult to generate specific research questions. Just trying to think them up can easily lead to mental blocks. Good ideas often come from reading, discussing, explaining (and best of all, teaching) what someone else is doing. Group discussions can be fertile breeding grounds for new ideas. Read current research papers in areas that interest you, force yourself to present and explain them to others, and ideas will strike you. In my experience it's not the authors' suggestions for future research that spawn the best questions: those suggestions are ones the authors themselves haven't been able (or bothered) to pursue successfully. People who write research papers generally know far more about what they are doing than the reader, and problems that they identify but leave unsolved may well be really tough! It's better to capitalize on your more detached position to escape from the author's mind-set and think more laterally about what he's working on, rather than joining him in the tunnel of his vision and identifying open issues through his eyes.

2.2. Relating ideas to a theme

Do not base a research proposal on just one solitary idea. Strive to give your research some breadth of scope and long-term continuity, without appearing to spread yourself too thinly. This is not easy to achieve, but merits serious effort. As months stretch into years and years into decades, your results should build up and strengthen each other so that real progress can be perceived towards answering significant and difficult questions.

An alternative research strategy is more opportunistic: to identify problems that others have formulated but failed to solve properly, and jump in with a new technique of which they are unaware and show how it can be applied. This kind of predatory strategy is often adopted by those who have special knowledge of — or an obsession with! — a particular viewpoint or tool. One danger is that to a man with a hammer, everything looks like a nail: you may be blind to the inappropriateness of your pet methodology for many of the applications you investigate. Another is that while good and plentiful results may be obtained quite quickly, over the long term the research program as a whole may take on a scrappy, uncoordinated, character. It seems better to focus your long-term efforts on particular kinds of problem than on particular kinds of solution.

2.3. Safe versus risky research

Do not put all your eggs into one basket by describing only one research idea. By its very nature, it is hard to plan research, and any avenue — no matter how good it seems — may turn out to be sterile, infeasible, or simply incorrect. On the other hand, beware of promising to work on too many things, for your proposal will be criticized as being "unfocused." Reviews of proposals sometimes state explicitly that the evaluation would have been higher if fewer ideas had been included. You can spoil a good proposal by adding more to it.

Propose a mix of questions to work on; some short term and obviously answerable, others longer term, more risky, but potentially more valuable. It is important to take chances in research, and equally important to be aware of the risks being taken. Kuhn (1970) defines "normal science" as research firmly based upon one or more past scientific achievements, achievements that are acknowledged by the scientific community to supply the foundation for further practice. He contrasts this with "scientific revolutions" that question and re-structure established practice: "non-cumulative developmental episodes in which an older paradigm is replaced in whole or in part by an incompatible new one." Kuhn's distinction, which is designed for a grand scale (like Copernicus's or Einstein's revolutions in physics), also applies in miniature at the level of the individual researcher: safe versus risky research. Be aware of this distinction and propose work on different levels.

2.4. Evaluating research ideas

You have to evaluate your own ideas, assess their strengths and weaknesses, sharpen them, and present them in the most favorable light.

When you specify a goal, how will you know if you reach it? Of course, you may not expect to attain your goals, but if by chance you achieve complete success you ought to be able to tell that you have done so! Many research proposals specify goals that are so vague they could never be reached (or already have been — sometimes it's difficult to tell). It is essential to formulate goals sufficiently precisely that it will be possible to determine when they have been reached, and (if it's not completely obvious) you must explain how you will know. Goals that are stated in a way that makes it difficult to decide if they have already been achieved, or ones that are clearly completely out of reach, will destroy the credibility of any proposal.

Are your goals worthwhile, and why? The onus is on you to convince your reviewers that, if you are successful, you will have accomplished something worth doing. Of course, you might fail. But if you do succeed it is reasonable to ask what contributions will have been made to scientific knowledge (i.e. results that others can build on) or to practice (i.e. general techniques that others can use too). If you intend to prove a theorem that no-one cares about, or tackle a particular application in a way that does not shed light on others too, then research funding will be much harder to obtain (although in the latter case the application may be sufficiently useful in its own right that you can convince someone to pay for it as a development project).

Have you identified a rational approach — or, better, a few possible approaches — to tackling your chosen problem? Of course, it's very difficult to plan research, and that makes many people cavil at the very idea of a research proposal. But it is certainly reasonable to expect you to have some idea how to start. Obviously you should be able to say what you will do in the first few months. And you must plan something more than just "waiting for inspiration" or even "reading about the problem (and waiting for inspiration)"! Since research is evidently unpredictable and difficult to plan, have several different lines of attack in case some go wrong or do not pan out.

3. The research proposal

Given that you have the ideas, how do you describe them and make them sound worth funding? You should consider the impact of your presentation on a busy researcher — like yourself, though perhaps more experienced. You are describing your ideas to a colleague, not a business promoter. Your basic problem, as pointed out by Bundy (1988), is threefold: to convince the selection committee that:

- you have identified a well-formulated goal;
- attaining this goal is a significant contribution to computer science;
- you have a good chance of reaching it with the resources requested.

3.1. Describing your ideas

Your proposal will be evaluated by experienced, and probably sympathetic, researchers. They've been through it all themselves. They understand the difficulty of doing research and how hard it is to write a proposal. They realize that research is difficult to plan. They don't expect to be able to glean every last detail about what you want to do just by reading the proposal. But they can tell a lot about you, and the way you think, from your writing. They expect you to have thought pretty hard about your ideas, and to have worked conscientiously to explain and present them as clearly and straightforwardly as possible. They want to give you a chance, but they must justify it to themselves (and to others too). It's up to you to provide the evidence for a positive decision.

Don't make your research description a sales brochure. The kind of people who evaluate it will probably react negatively to salesmanship. On the other hand, you must make it clear that what you propose to do is worthwhile and has a good chance of success.

Acknowledge difficulties honestly. Don't try to pull the wool over the reader's eyes — he or she is probably pretty bright. If there are snags or potential problems, say so: reviewers will be impressed by your candor. If the difficulties are ones they haven't thought of, they may be impressed by your intelligence too. It is only reasonable to assume that you have thought through your proposal more thoroughly than the reviewers have; consequently if they see problems that you don't seem to have noticed then they will be less than impressed with your efforts. It would reflect badly on your proposal if you were to describe obstacles that seem completely insurmountable, but you presumably won't be proposing work that you judge to be quite infeasible. You cannot really lose by being honest about the problems you expect to encounter.

3.2. The researcher

As well as having good ideas, you must explain why you're fully — perhaps uniquely — qualified to carry them out. Of course, since they're *your* ideas, you automatically have a head start over others.

You must know the background to the work, the relevant literature, and what others have done. Your proposal should contain a section that reviews prior work. Space will not permit a comprehensive literature survey, and you will be unable to include many references. That makes it all the more important to select judiciously, thereby demonstrating that you have solid knowledge of the field, and the ability and good taste to make the very best use of limited space. Do not be overly introverted: mention other work besides your own. It gives a bad impression to have all (or even most) references to yourself or to a closed circle of collaborators. Avoid being involved in a small clique of researchers who publish in the same places and whose results are referred to only by one another.

For a senior researcher, the "track record" of work in the area will obviously play an important role in the evaluation of the proposal. Do not waste space by listing your own papers twice, once in the reference list and again in the personal data form or résumé. Invent a way to cross-reference from the proposal to the personal data (e.g. by numbering entries in your publication list and using letters to identify other references in the proposal). If you do not have an extensive track record, do not fret — your proposal will be judged relative to others at similar stages in their career. Everyone has to begin somewhere; the people who evaluate your proposal know that.

There is not much you can do to boost your track record, other than presenting your accomplishments fairly and accurately (see section 4). However, another most important source of information concerning whether you are the right person for the job is the understanding and insight you display when presenting and discussing the research in the proposal. The payoff for explaining your ideas clearly, eloquently, insightfully and candidly cannot be stressed too strongly.

3.3. Structure of a proposal

Any proposal should review the context of the research, articulate the goals that will be pursued, summarize relevant prior work, describe a research plan, and give some indication of why the research is useful. Sometimes it is necessary to include a progress report on already-completed research as well.

The background should be brief and set the context for the proposal in terms of an overall research theme. The goals should project a fabric of interwoven ideas, augmenting and contributing to each other, with a mix of shorter- and longer-term, safe and risky, research, so that even if some ideas turn sour plenty will get done. One useful technique is to break down an overall goal into several interacting sub-goals or objectives — but beware of proposing too much.

For the research plan, you should at least know how you're going to start out and have some ideas for future options. Don't schedule research too firmly or too far into the future; that's unrealistic. Be prepared to describe alternative scenarios for the later stages, which hinge on how the early research turns out. Look at the problem from different points of view (theory, simulation, experimental implementations, human behavior, ...) to make it clear that you have in mind a rich variety of approaches, and the personal resources to carry them out.

Be very mindful of the need to *evaluate* your ideas, not just develop and implement them. If successful, what will be the effect of the research — how will others be able to build on the results? Will you contribute to the advancement of science, or merely develop a wonderful "look Ma, no hands" system that leaves others no better off? Sometimes such systems leave others worse off; they cannot replicate or follow up on your results and therefore cannot credibly pursue that line of research themselves.

3.4. The progress report

If you have previously been funded, you must summarize progress under earlier grants. What specific contributions have been made, where have they been published, who has taken them up, applied them, or developed them further? What students have been trained, what papers have *they* written, who has hired them? If you cannot demonstrate that you have made good use of a previous grant, your chances of getting a fresh one will clearly be diminished.

Publication delays may mean that your recent work has not yet appeared in print. Some papers based on work prior to a previously-held grant may have appeared; while this is a good sign it should not be confused with research progress stemming from the grant itself. Fortunately, your proposal will be evaluated by experienced researchers who understand the publication business and get frustrated by publication delays themselves.

3.5. Preparing the proposal

The people who evaluate your proposal are *busy*, even overloaded. Moreover, they are volunteers! If you can't be bothered to take the trouble to present your ideas clearly, why should they bother to read them carefully? There's a lot in it for you, much less for them. Of course, we all know that preparing research proposals is a nuisance, but reading them (by the dozen or hundred!) can be far worse. Readers will react very negatively to any signs of sloppiness in either thinking (fuzzy goals, inadequate background, unacknowledged problems ...) or presentation (poor proof-reading, spelling errors, infelicitous formatting, incomplete references ...). If you aren't sufficiently motivated or excited by your ideas to spend time honing the content and presentation of your proposal, you can't expect a sympathetic hearing from whoever is obliged to evaluate it.

Just because the reviewers are busy does not mean they will look favorably on a superficial or "popularized" proposal. Make sure there is plenty of technical content for them to pick up on. If the proposed research is highly technical, do not shy away from reflecting the technicalities in the proposal. There is nothing wrong with including a few equations if necessary, even diagrams (though be careful, especially with the latter, to ensure good use of space). Although your proposal must of necessity be brief, do not make it anaemic.

Have others read your proposal before submitting it. Encourage them to be critical, to emulate a tough reviewer, to pick out holes and ambiguities, to misunderstand where at all possible — in short, to look for ways to dislike the proposal. Probably the actual reviewers will be much more sympathetic, but you should prepare the proposal to withstand critical onslaught.

Proposals are restricted to a certain number of pages. You don't have to cover them all, but a clear exposition of complex ideas takes a certain amount of writing and most successful proposals occupy the majority of the allotted space. Don't buck the system by using a tiny typeface. Prepare the proposal in a straightforward way that won't upset the reader. It is better to get the bulk of your message across properly than to try to communicate the whole thing in detail and fail completely! Don't try to cheat by sending in more than the maximum number of pages: the proposal will be truncated before it even reaches the reviewer and the really important parts may be lost. Think of it as an exercise: part of the test is seeing how effectively you can work within specified constraints.

4. The personal data form

Along with the research proposal you will have to submit a personal data form giving information about your qualifications, the positions you have held (list them in reverse chronological order), the number of students you have supervised (specify co-supervision; divide Master's students into coursework and thesis students if applicable), your publication list, and other information. Make sure you document industrial and consulting work, along with any other "technology transfer" activity. Consider showing thesis titles and other publications by students under your supervision, listing your graduate students by name, summarizing your refereeing activity, your published reviews, and so on. What you decide to include reflects your priorities and general professionalism; it will be used by the reviewer to build a picture of you and your work.

4.1. The publication list

This is perhaps the most important part of the personal data form and you should take great care in preparing it. Gather together under separate headings papers in refereed journals (clearly indicating their "accepted" or "published" status), papers in refereed conference proceedings, other refereed items like book chapters, books, non-refereed articles, and so on. Make sure you have complete references to your papers and check that you give them the correct titles (it's surprising how many people don't!).

It is essential to be scrupulously honest when preparing the publication list. Reviewers react very negatively to any suspicion of cheating. Make sure you know for certain which of your publications are refereed. Journals, even high-profile ones, for which papers are accepted on the judgement of the editors alone are not refereed — even if one or two members of an editorial board are consulted too.

Avoid duplication in your publication list. If a conference paper was subsequently published as a book chapter, for example, choose one section in which to include it and note with that entry that it also appeared elsewhere. In general, if it's a reprint or a revision of an earlier paper, say so, and only list it once (you do gain credit from the fact that someone evidently thought it was worth reprinting!). Do not write different papers with the same (or very similar) titles.

Submitted papers should be collected together and clearly identified as such. People disagree on whether you should specify the journals to which your papers have been submitted. The argument in favor is that it gives readers a chance to judge whether you are submitting your work to appropriate places. On the other hand it might be interpreted as an attempt to glorify yourself by association with these revered organs. Moreover if your paper is not accepted you risk exposing your failure: it is surprising how much reviewers remember from one grant application to the next!

Never succumb to a temptation to mislead reviewers on the status of submitted papers — it's quite possible that someone will check with the editor of the journal and discover deception (it happens). If a paper has survived one round of refereeing and been re-submitted for a second, say so. If it has been accepted subject to minor corrections and approval by the editor, say so, giving the date of acceptance. If in doubt, spell it out.

These remarks are intended as guidelines rather than rules, and in practice there is some latitude in interpreting them. Some people prefer to list duplicate papers under each category in which they have been published, which is permissible so long as they are clearly cross-referenced. The refereed or non-refereed status of papers is sometimes not clear-cut, particularly in the case of invited papers — and ultimately, of course, it is the quality of the material that counts, not where it appears. The most important thing is to be open and honest about the status of your work. *If you are suspected of misrepresentation, your application will suffer and so will your reputation.*

4.2. Additional material

You may have an opportunity to submit additional material, such as preprints or reprints, to support your application. Unfortunately, reviewers are often forced to guess the quality of a paper from the journal or conference in which it appears, but if you can submit actual papers this provides a welcome opportunity to evaluate the research itself. Be sure to select reasonably recent work, and make it your best work! Do not include papers just because they have been published in prestigious journals. It may be better to choose good papers that have appeared in obscure places, or have not yet been published, as the reviewer will otherwise be quite unable to evaluate this work.

4.3. Updates

You may have an opportunity to submit an update to your personal data form after the application has been submitted. If you want to really annoy the reviewers, make a few minor changes to your publication list, re-sort it in a completely different order, and re-format it: then they will have to go painstakingly through the two lists to spot the differences (and determine that they are insignificant)! Although document preparation technology may make it easier for you to re-format and re-print the entire list, it will be far easier for the reader if you simply prepare a list of the differences, specifying clearly what has been changed or added, and how it should now read.

5. How a selection committee works

It helps to know a little about how a grant selection committee works. The committee has twelve or so members, each of whom reads every application in advance. Two or three members who are especially knowledgeable about the relevant research area are specifically assigned to each proposal as "internal reviewers" to evaluate it thoroughly, prepare a recommendation, and present it to the committee. The natural tendency is for internal reviewers to champion their applications where merited, and the other committee members will serve as a critical sounding-board for the presentation. The meeting proceeds quickly. Internal reviewers say a little about your application, highlighting your credentials, what you propose to do, their evaluation, and their recommendation. If your application is any good they will be on your side, trying to persuade the other committee members of the virtues of your case. You should strive to make it easy for them! Re-read your application and imagine someone having to defend it on your behalf in the space of a few minutes. Obviously you must highlight salient points in the summary: goals, prior achievements, objectives, research plan, evaluation methodology.

Meanwhile, as your representatives present your case, the other members of the committee are leafing through the application (probably the personal data form), trying to assess the case and whether they can agree with the recommendations or not. They have studied it before, of course, but there may be hundreds of other applications and memories will need refreshing. Table 1, adapted from Bundy (1988), summarizes common reasons why proposals are rejected — bear these in mind as you prepare your proposal. There might be disagreement between the internal reviewers or with another committee member — an argument! As the discussion proceeds, the rest of the committee is silently scanning your application, listening, and thinking about it. Just imagine the impact of a poorly-prepared, scrappy proposal, and contrast it with the effect of a beautiful, tastefully-arranged document.

6. Refereeing grant applications

Selection committees depend heavily on timely and careful reviews by outside members of the research community. Each application is sent to several external referees for evaluation. Some are suggested by the applicant, others by the committee. The responses are made available to all members and referred to frequently in the committee's discussion. Indeed, in the case where no committee member has direct research expertise in the area of the application — or when the only member who has cannot contribute because of a conflict of interest — external reviews may form the primary basis for evaluation.

Refereeing other people's applications is widely perceived as a time-consuming chore, although it can be very interesting. But just think how much more onerous it is to be on the committee itself, charged with reading hundreds of applications instead of just a handful! Ultimately it is in our discipline's interests to have the fairest possible funding decisions, and conscientious reviews play a crucial role in this. For example, NSERC evaluates the functioning of the computer and information science committee, and the computer science community at large, by the response rate to review requests: this is the kind of thing that helps whenever the committee makes requests for a larger slice of the cake. If you care about funding for computer science, you should feel obliged to contribute your share to the refereeing process.

- It is not clear what question is being addressed by the proposal.
- It is not clear what the outcome of the research might be, or what would constitute success or failure.
- The question being addressed is woolly or ill-formed.
- It is not clear why the question is worth addressing.
- The proposal is just a routine application of known techniques.
- Industry ought to be doing it instead.
- There is no evidence that the proposer has new ideas that make it possible to succeed where others have failed.
- A new idea is claimed but insufficient details are given to judge whether it looks promising.
- The proposer seems unaware of related research.
- The proposed research has already been done (or appears to have been done).
- The proposer seems to be attempting too much for the funding requested and the time-scale envisaged.
- The proposal is too expensive for the probable gain.

Table 1. Some reasons for rejecting a research proposal (adapted from Bundy, 1988).

It is important to prepare reviews thoughtfully and to the best of your ability. Unqualified praise gives the impression that you are trying to do the applicant a favor; unqualified criticism that you have a biased view. In any case it is helpful for you to summarize your previous knowledge of the applicant's work and your personal acquaintance of him or her, if any. One-line reviews give the impression that you haven't taken time to reflect upon the proposal or evaluate it properly. On the other hand, no-one wants to read a review that is longer than the proposal itself (yes, it does happen!). The best reviewers evaluate proposals carefully and summarize the evaluation fairly, mentioning both positive and negative aspects and weighing the evidence for and against funding. Writing good reviews is just another aspect of your professionalism: it will be noticed and will enhance your reputation.

Acknowledgements

I am grateful to Rick Bunt, Brian Gaines, Saul Greenberg, Carl Hamacher and Helmut Jurgensen for making valuable comments on a draft of this article.

References

Bundy, A. (1988) "How to get an SERC grant," *AISB Quarterly 65*: 7–9.

Kuhn, T.S. (1970) *The structure of scientific revolutions.* University of Chicago Press, second edition.

KR'91 — Call for Papers

Second International Conference on Principles of Knowledge Representation and Reasoning

Royal Sonesta Hotel Cambridge, Cambridge, Massachusets
April 22-25, 1991

The idea of explicit representations of knowledge, manipulated by general-purpose inference algorithms, underlies much work in Artificial Intelligence, from natural language to expert systems. A growing number of researchers are interested in the principles governing systems based on this idea. This conference will bring together these researchers in a more intimate setting than that of the general AI conferences, and authors will have the opportunity to give presentations of adequate length to present substantial results.

Topics of Interest

Submissions are encouraged in a least the following area:
Knowledge representation formalisms such as the logics of knowledge and belief, nonmonotonic logics, temporal and spacial logics, taxonomic logics, and logics of uncertainty and evidence;
Generic ontolgies for describing time and space, resources, causality, and constraints; Reasoning methods for knowledge representation systems such as abduction, induction, belief management and revision, deduction, analogical reasoning, learning, planning and plan analysis, diagnosis, inheritance and classification.

Submissions for Abstracts

Send Five (5) copies of an extended abstract (maximum 8 pages) by October 22, 1990 to one of the program chairs listed here:

Erik Sandewall, Computer & Info. Science
Linkoeping University
S-58183 Linkoeping, Sweden
+46 1328 1408 email: ejs@IDA,LiU.SE

Richard Fikes, Price Waterhouse Technology Centre
68 Willow Road
Menlo Park, CA 94025
(415) 688-6684 email: pwtc!fikes@labrea,stanford.edu

For general information about the conference, please contact the conference chair;
James Allen, Computer Science
University of Rochester
Rochester, NY 14627
(716) 275-5288 email: james@cs.rochester.edu

Important Dates

Submissions receipt date: October 22, 1990
Author notification date: December 3, 1990
Camera-ready copy due: February 1, 1991

Mainframe Knowledge-Based Systems

by Bruce D. Scott

Perhaps the primary lesson to be learned about mainframe Knowledge-Based Systems (KBS) is that to realize the full potential of the technology on a new platform, the model of the standalone, consultative expert system developed on workstations requires significant rethinking. To date, the biggest payback for mainframe-based expert systems has been on systems highly-integrated into the mainframe environment, or in many cases, on systems completely embedded within other applications, with no user interface. Standalone applications, not making use of databases, electronic mail, on-line utilities, and the like, generally only make sense in a mainframe environment when the benefits of shared code outweigh other considerations.

From our perspective, the future for mainframe KBS is as an integrated element of mainstream data processing practices, and not as an office automation-type oddity standing somewhere on the fringes. Monolithic helpdesk-type applications have their place but many mainframe shops working with KBS technology are discovering greater paybacks in batch KBS, or embedded KBS called as sub-routines in larger transaction processing systems such as IMS or CICS.

This, of course, means that mainframe KBS developers must address certain issues not ordinarily of paramount concern to workstation-based developers. Systems that process thousands of transactions/consultations daily must perform completely predictably, without exception. Consultations which update production databases must terminate successfully and correctly, since failures may not be detected until considerable damage has been done. Mainframe KBS must be robust; not subject to frailty for which expert systems are known, and must support multiple concurrent users. Efficient execution is also a major concern, especially for service bureaus. Few mainframe cycles go unaccounted for.

Finally, mainframe KBS software must integrate seamlessly with traditional DP languages such as PL/l and COBOL, and industry standard databases such as IMS and DB2. Survivors of the data processing wars have seen many oversold technologies come and go, and are loath to adopt anything untried especially if it involves abandoning proven practices and technologies.

Of course, programming "smart" DP systems is nothing new. Many data processing veterans can point to systems developed in the past which utilize specialized domain-specific expertise. What is new in the mainframe environment is the emergence of AI-derived design and coding techniques which can be added to the data processing arsenal. A number of expert system "shells" are being marketed for the mainframe environment which support the use of such AI-derived techniques.

KBS technology is not about tools, however, but about techniques. Knowledge representation and inferencing techniques, such as frames, patterning matching, and chaining, are the central issues. Issues of importance to software vendors rather than system developers too often dominant the discussion of KBS. Successful applications are not the necessary consequence of development using KBS software.

Tools are of interest, but only in so far as they support the use of KBS techniques. Beyond that, one tool is superior to another according to the usual criteria of flexibility, ease-of-use, efficiency, ease-of-integration into the DP environment, and so on. Considering that no "off-the-rack" KBS tool can be all things to all people and that fielding more than one expensive KBS package creates headaches of its own, acquiring a single open architecture, enhanceable package is usually preferable. Open architecture is important since one can code and link in features missing from the original shell. Finding such a tool, however, still remains a problem since fully-featured, industrial-strength mainframe KBS software is rare and writing in LISP is simply not an option for most mainframe developers.

In regard to the skill set required to construct mainframe KBS, suffice it to say, a hybrid skill set is required of mainframe KBS practitioners, combining both traditional DP skills and Knowledge Engineering skills. Unfortunately mainframe KBS development often requires skills unfamiliar, if not alien, to most DP practitioners. Unstructured logic, iterative development, constant interaction with users and experts, and non-procedural coding are all features of KBS development that most DP shops have worked very hard to eliminate. On the other hand, trained Knowledge Engineers rarely have much exposure to the technical demands of mainframe application development.

Mainframe development also involves a slightly different mix of KBS skills than development in the workstation environment. Issues concerning system correctness and maintainability, automatic performance auditing, error diagnosis and recovery rise in importance since embedded or batch systems run more untended than traditional interactive expert systems. Conversely, issues of user interface design, consultation flow, and backward chaining decline in relative importance. Batch or embedded KBS often have all required inputs available up front so forward chaining rule bases alone can often guarantee the desired results.

Other major differences can be traced to the use of large volumes of on-line data. Mainframe KBS are often smaller than their workstation-based cousins, with rule counts, for example, often less than 100 rules. Much of what would be captured in a standalone knowledge base may already exist in mainframe databases. KBS applications which are designed to utilize such data can often provide excellent functionality while at the same time remaining small and inexpensive to run and maintain.

Knowledge bases are used in such applications to capture metaknowledge (knowledge about knowledge/data) rather than the usual combination of both data and knowledge. A data-intensive KBS must know how to interpret the data. That is to say, it must know under what conditions certain facts are relevant, how to cluster observations and by what criteria, how to map logical or symbolic meaning onto physical data base layout, and so on. With thoughtfully-designed knowledge bases, KBS maintenance, an issue of some concern, becomes merely an issue of data base maintenance.

Smart relational data base front ends need know nothing about specific data bases, merely about relational data bases in general. Generic data base query utilities, for example, dynamically create menus which include only the options which are relevant given

---

Bruce Scott is a Senior Technical Analyst, working in the capacity of Knowledge-Based systems group leader and Senior Knowledge Engineer since September, 1987 with British Columbia Systems Corporation.

what can be inferred about the user and his interests, as well as the actual content of the database. Menus become less cluttered, unproductive data base queries are reduced, and processing costs are cut.

Data-intensive KBS applications, processing large volumes of information, need not necessarily handle all cases, merely the cases that constitute the regular or commonplace situations. In addition to correctly dealing with the regular cases the system must also recognize exceptional or extraordinary instances, setting them aside for human consideration or flagging them as uncertain. Handling 70% of 10,000 cases relieves the human expert of an enormous burden.

As an integral part of mainstream data processing, mainframe KBS technology holds considerable promise as a productivity tool. But realizing this potential involves selectively adopting only the DP-compatible techniques and operating principles from traditional workstation-based expert systems, and creatively integrating them with established data processing technology and practices.

Report on the Second International Workshop on User Modeling

Honolulu, Hawaii, March 29 to April 1, 1990

Paul van Arragon
Dept. of Computer Science, University of Waterloo

Cet atelier fut un succès grâce a l'intérêt grandissant pour la modélisation d'utilisateur. Vingt-trois articles, ce qui représente environ la mortié des articles soumis, ont été acceptés et présentés au cours de l'atelier. Parmis ces articles, dix représentaient du travail fait au niveau du doctorat.

Chaque article fut allowé une période de présentation de vingt minutes suivie par une période de discussion également de vingt minutes. Toutes les discussions furent très animées. La majorité des articles faisaient états des progrès des différents projets de recherches. On y retrouvait une impressionante variété de projet. Voici quelques uns des sujets qui soulevèrent le plus d'intérêt: les outils de modélisation d'utilisadeur, la reconnaissance des plans et le génération du langage naturel.

Dans les paragraphes qui suivent vous trouverez un compterendu de certains articles présentés.

General Chair: Professor Wolfgang Wahlster, University of Saarbrücken, West Germany
Program and Local Arrangements Chair: David Chin, University of Hawaii, USA

Growing interest in the topic of user modeling made this workshop a success. Twenty-three papers, about half of those submitted, were accepted and presented. Ten of these represented Ph.D. work. Each paper was presented for twenty minutes with twenty minutes allotted for discussion. Discussion periods were lively without exception.

Most of the papers focussed on describing one's own current research. The topics covered a wide spectrum. "Hot" topics include user modeling tools, natural language generation, and plan recognition. Below we provide brief summaries of work presented.

1. Psychological Foundations

Martha Crosby and Jan Stelovsky of the University of Hawaii monitored eye movement to explore how computer science students were studying Pascal programs. Expert students could be distinguished from novices.

Paul McKevitt of New Mexico State University used Wizard-of-Oz experiments to study what types of questioning sequences are common when students interact with computers. Knowing these sequences is useful for resolving anaphoric references.

Ira Haimowitz of the MIT Laboratory for Computer Science used user models to generate empathetic responses. The work is applied to medical domains where patients need to be reassured that they can leave the hospital shortly and that treatment will not be too painful.

2. Student Modeling

Gordon McCalla and Jim Greer of the University of Saskatchewan studied how to maintain an accurate model of a student whose knowledge is constantly changing. A representation of knowledge that permits granularity shifts is useful.

3. Plan Recognition

Douglas Appelt and Martha Pollack of SRI International developed a formalism called *weighted abduction.* It permits the comparison of plans when more than one plan is possible.

Alex Quilici of UCLA used domain-independent knowledge about plans to model user's incorrect beliefs about plans. This permits responding to misconceptions.

Robin Cohen, Fei Song, Bruce Spencer, and Peter van Beek of the University of Waterloo extended plan recognition to recognize temporal constraints and admit novel information .

Dekai Wu of UC Berkeley acquired user plans by querying the user during dialogue, and finding the plan that maximizes utility.
Sandra Carberry of the University of Delaware used preferential rules based on Dempster-Schafer theory to choose among competing plans.

Robert London of Cimflex Teknowledge, USA developed a student modeler, IMAGE, that supports multiple approaches to instruction by anticipating the student's plans.

Rhonda Eller and Sandra Carberry of the University of Delaware used dialogue to detect discrepancies between the model of a user's plans and the actual plan.

4. Theoretical Issues

Paul van Arragon of the University of Waterloo developed theoretical tools for reasoning about a user's default assumptions. In the tools, both the system and the user reason by default.

Alfred Kobsa of the University of Saarbrücken, West Germany, developed BGP-MS, a workbench for developing user models. A rich representation language exists for conceptual knowledge, and an interface exists that enhances many activities, including customizing stereotype management .

Afzal Ballim of the Inst. Dalle Molle pour les Etudes Semantiques et Cognitives, Switzerland and Yorick Wilks of New Mexico State University dynamically ascribed nested beliefs using complex stereotypes.

Franz Schmalhofer and Otto Kuehn of DFKI, West Germany used mental models to build how-to-do-it knowledge.

Judy Kay of University of Sydney developed a toolkit that allows the user to modify his or her own model.

5. Natural Language Dialog

Susan Brennan of Stanford University compared human/human dialogue with human/computer dialogue, and found many similarities and differences. One difference is that in human/computer dialogue, there is less grounding: establishing that understanding took place.

Harry Bunt of Tilburg University, the Netherlands, developed a formalism for incrementally modeling beliefs and intentions. The formalism makes use of partial models of a formal language.

Robert Kass of EDS, USA acquired user models using thirteen default rules based on features of human conversation. The resulting system acquires user models during conversation .

6. Natural Language Generation

Ursula Wolz of Columbia University showed how the user model of the question-answering system GENIE influences generated responses.

Ingrid Zukerman of Monash University, Australia developed mechanisms to anticipate the effect of given utterances on listeners, and to avoid adverse effects.

Mark Maybury of Rome Air Development Center, USA developed a model to plan utterances by reasoning about the rhetorical structure and pragmatic function of the text, guided by the user model.

Margaret Sarner and Sandra Carberry of the University of Delaware generated definitions tailored to a user's needs in a task-oriented dialogue by taking several aspects of the user into account.

7. Canadian Content

As listed above, three talks issued from Canadian universities. A style that emerged was for one speaker to present the paper, to try to incite controversy, and then to let the other author field questions. Jim Greer did this to Gord McCalla, and Peter van Beek did this to Robin Cohen.

Jim Greer also distinguished himself by an attempt at humour. After Robin and Peter's talk, which discussed tense in natural language, Jim piped up, "All this talk of tense is inappropriate. This is Hawaii; we should be hanging loose." A few people laughed politely.

8. Hawaiian Content

Having the workshop in Hawaii seemed to be a wise choice. Although inconvenient for the European attendees, the good weather helped draw people to the workshop.

No banquet was prearranged for the workshop. Several participants suggested that we go en masse to a Hawaiian luau to see Hawaiian dancers, play Polynesian games, drink Blue Hawaiians, and eat Hawaiian food. David Chin, who lives in Hawaii and planned local arrangements, tried to discourage this idea on the grounds that it was touristy kitsch, but he was overruled almost unanimously.

A good time was had by all, although at times our group seemed unimpressionable. Our Polynesian guide was frustrated when trying to make us laugh, because we refused to be amused. The lowest moment came when, between dances at the luau, the overweight Polynesian host was welcoming all the groups present. Unfortunately, he misread his notes and consequently welcomed the "UH workshop" instead of the "UM workshop". He chided us for our unenthusiastic response, and had already introduced two or three other, more-enthusiastic groups before we successfully resolved his incorrect reference.

9. Future User Modeling Events

User modeling research is gaining momentum. As announced in Canadian Artificial Intelligence, April 1990, a new user modeling journal (UMUAI) is being created. Selected papers from the workshop will appear there. Also available are a limited number of reprints of the workshop papers. (To obtain one, ask Alfred Kobsa at "kobsa@cs.uni-sb.de".)

Further user modeling workshops are being planned to take place in a castle in West Germany in 1992, and at IJCAI in Australia, 1991 and France, 1993. There may also be other user-modeling events at IJCAI 1993, such as a tutorial.

General Chairman
Moonis Ali
Univ. of Tenn. Space Inst.

Program Chairman
Jim Bezdek
University of West Florida

Program Co-Chairs
Ken Ford
University of West Florida
Takeshi Yamakawa
Kyushu Inst. of Tech.

Program Committee
Jack Adams-Webber
Brock University
Jeff Bradshaw
Boeing Computer Services
Bruce Buchanan
University of Pittsburgh
Eric Dietrich
SUNY - Binghamton
D. Dubois / H. Prade
Paul Sabatier Univ.
Jim Etheredge
Cal Poly State Univ. SLO
Doug Fisher
Vanderbilt University
Tom Glynn
Glynn Scientific Inc.
Avelino Gonzalez
Univ. of Central Florida
Edward Grant
Turing Institute
Mehdi Harandi
University of Illinois
Pat Hayes
Xerox PARC
Michael Huhns
MCC Corporation
Victor Law
Tulane University
Bob Lea
NASA/JSC
Tod Levitt
Advanced Decision Systems
Manton Matthews
University of South Carolina
Penny Nii
Stanford University
Arnold Norman
Rockwell International
Margaret Nowak
Harris Corporation
Judea Pearl
UCLA
Fred Petry
Tulane University
John Roach
Virginia Tech. University
Erik Sandewall
Univ. of Linköping
Mohan Trivedi
University of Tennessee
Jay Weber
Lockheed AI Center
Bernie Widrow
Stanford University
Ronald Yager
Iona College
Hans Zimmermann
Aachen Inst. Tech.

CALL FOR PAPERS

The Fourth International Conference on Industrial & Engineering Applications of Artificial Intelligence and Expert Systems

IEA/AIE-91

June 2-5, 1991 Waiohai Hotel Kauai, Hawaii

Sponsored by:

The University of Tennessee Space Institute (UTSI)

In Cooperation with:

American Association for Artificial Intelligence
Association for Computing Machinery/SIGART
Canadian Society for Computational Studies of Intelligence
European Coordinating Committee for AI
IEEE Computer Society
International Association of Knowledge Engineers
International Neural Network Society
Japanese Society of Artificial Intelligence

This conference continues its tradition of emphasizing applications of artificial intelligence and expert/knowledge-based systems to engineering and industrial problems. Topics of interest include but are not limited to:

- **Pattern Recognition**
- **Neural Networks**
- **Machine Learning**
- **Machine Vision**
- **Computer Aided Manufacturing**
- **Computer Aided Design**
- **Expert & Diagnostic Systems**
- **Intelligent Databases**
- **Intelligent Interfaces**
- **Intelligent Tutoring**
- **Knowledge Representation**
- **Knowledge Acquisition**
- **Reasoning Under Uncertainty**
- **Model-Based & Qualitive Reasoning**
- **Natural Language Processing**
- **Robotics**
- **Distributed AI Architectures**
- **Planning/Scheduling**
- **Sensor Fusion**
- **Fuzzy Logic**

Please indicate the area of submission as shown above, and submit by October 1, 1990 three copies of an extended abstract (6-8 single spaced pages) to the Program Chair. Authors will be notified of the status of their manuscript by December 15, 1990. Final copies of papers will be due for inclusion in the conference proceedings by March 1, 1991.

Dr. Jim Bezdek, Program Chair
Division of Computer Science
University of West Florida
Pensacola, FL 32514

Tel. (904) 474-2784; Fax (904) 474-3129
E-mail: JBEZDEK@UWF.BITNET

Dr. Moonis Ali, General Chair
Computer Science Dept.
Univ. of Tennessee Space Inst.
Tullahoma, TN 37388

Tel. (615) 455-0631 ext.236
E-mail: ALIF@UTSIV1.BITNET

General information and registration materials may be obtained from:
Ms. Sandy Shankle,
UTSI; Tullahoma, TN 37388-8897.
Tel. (615) 455-0631, ext. 278.

The proceedings will be published by ACM and will be available at the conference.
Proceedings of earlier conferences are available -- contact:
Ms. Nancy Wise at (615) 455-0631, ext. 236.

BOOK REVIEWS

Edited by Graeme Hirst

Readings in qualitative reasoning about physical systems
Daniel S. Weld and Johan de Kleer (editors)
[University of Washington and Xerox Palo Alto Research Center]
San Mateo, CA: Morgan Kaufmann Publishers, 1990, x+720 pp (The Morgan Kaufmann series in representation and reasoning) Paperbound, ISBN 1-55860-095-7, us $34.95
(Distributed in Canada by John Wiley and Sons Canada Ltd)

Reviewed by
Howard J. Hamilton
Simon Fraser University

Capsule review: *Readings in qualitative reasoning about physical systems* (hereafter *Readings*) is a good, reasonably complete collection of research papers from 1975 to 1988 dealing with qualitative reasoning (also known as "qualitative physics"). It is suitable for reference or for background reading for people new to the area. If this description matches your needs, go ahead and buy *Readings* because it is a worthwhile investment compared to seeking out the many papers individually.

Introduction: *Readings* is another in Morgan Kaufmann's series of collections of articles. Each book in the series is a collection of noteworthy articles on an area of AI research, selected and introduced by one or more authorities on the area. The articles are typically photo-reproductions of the versions which appeared in journals, proceedings, or books, although sometimes an updated version is included instead.

De Kleer, the influential founder of the qualitative reasoning research area, is certainly an appropriate choice to edit the collection. He and Weld have done a fair bit of work in putting together the collection. Their selection of articles is good, their introductory comments are interesting, and they include a bibliography and name and subject indexes.

History: According to Weld and de Kleer, the goal of the research area is to build systems that "...[are] capable of reasoning about the physical world much as we ourselves, as engineers and scientists, do" (p. 1). Instead of an introductory overview of the field, the editors supply some historical notes. Some of the phrasing in this introduction is awkward (it begins with "Dan and I", but later says "we ... present Johan's reasons for becoming interested in qualitative physics"), but de Kleer's history of the idea of qualitative reasoning is interesting enough to compensate for any awkwardness. In 1974, de Kleer participated in the "confusion seminar," run by Seymour Papert at MIT. The purpose of the seminar was to analyze thinking by studying confounding puzzles about commonsense topics such as bouncing balls, roller coasters, etc. The lesson that de Kleer learned was that most problems could be solved by simple qualitative reasoning with a few simple equations. But the conventional mathematical formulation of physics was not helpful for these problems for one of three reasons: (l) the equations were intractable; (2) people were not convinced by the answers unless an intuitive explanation was supplied; or (3) the wrong equations were chosen. Like Hayes in his "naive physics" manifestos (the second of which is included in *Readings*), de Kleer realized the importance of commonsense knowledge to physics; he says, the so-called physics knowledge of kinematics and Newton's laws comprise[s] only a small fraction of the knowledge needed to solve problems. Most of the knowledge is "pre-physics, and considerable effort is required to codify it" (pp. 2-3). In 1984, a special issue of the journal *Artificial Intelligence* dealt with qualitative reasoning about physical systems. A flood of papers on the subject followed at AAAI and IJCAI conferences. As yet, the flood has not abated.

Pause for reflection: The appearance of *Readings* provides a good opportunity for us to reflect on the achievements and failings of research into qualitative reasoning, just as editing the collection provided the same opportunity to Weld and de Kleer. Research on qualitative reasoning has not achieved or even focused on its original goals, de Kleer and Weld note with disappointment. They claim that much research has "built bridges over dry land" (p. 7) by investigating minor issues for their own sake, rather than applying the ideas to tasks. Certainly, few of the papers in *Readings* include applications, but the editors resist pointing fingers.

One issue worth reflecting on is whether qualitative reasoning techniques have been created using an appropriate method of formalizing commonsense. A *knowledge-oriented* method formalizes commonsense knowledge about liquids, etc., without regard for how the knowledge will be used; a *reasoning-oriented* method formalizes commonsense reasoning without regard for the structure of the knowledge that is represented. Lenat's CYC project (Lenat *et al*, 1986) and Hayes use knowledge-oriented methods. In his manifestos, Hayes advocated encoding commonsense knowledge in first-order logic, but Weld and de Kleer note that much of Hayes' advice about how to get the enterprise [of encoding naive physics] done was ignored "(and in our minds correctly so)" (p.10). Most qualitative reasoning research has used a reasoning-oriented method of formalizing: a *qualitative model*, a simple, discrete representation of a physical system is designed by hand, and then the researchers argue that by applying constraint propagation to this model an approximation of human qualitative reasoning results. Preoccupation with the formal properties of such reasoning systems may be the type of minor issue that Weld and de Kleer dislike. Weld and de Kleer also dislike the cost of encoding the physical theories painstakingly by hand and regret that each researcher has individualized representation and reasoning methods. (The latter problem dates at least from the 1984 *AI* journal special issue, where editor Daniel Bobrow states "the planning of this volume ... caused some convergence in the notation used by the authors, though not nearly enough for my taste".) Weld and de Kleer end their preface by saying:

> Perhaps qualitative physics should begin a CYC-like project to develop a common language for describing a physical world to be used throughout the qualitative physics research

Call for Participation

Twelfth International Joint Conference
August 24 - 30, 1991. Sydney. Australi

The biennial IJCAI conferences are the major forums for the international scientific exchange and presentation of AI research. The next IJCAI conference will be held in Sydney. Australia, 24-30 August 1991. IJAI-91 is sponsored by the International Joint Conferences on Artificial Intelligence. Inc. (IJCAII), and co-sponsored by the National Committee on Artificial Intelligence and Expert Systems of the Australian Computer Society.

The conference technical program will include workshops, tutorials, panels and invited talks. as well as tracks for paper and videotape presentations.

Paper Track:

Submission Requirements and Guidelines Topics of Interests.

Submissions are invited on substantial, original, and previously unpublished research in all aspects of AI, including, but not limited to:

- Architectures and languages for AI (e.g. hardware and software for building AU systems, real time and distributed AI)
- Automated reasoning (e.g. therom proving, automatic programming, planning and reasoning about action, search, truth maintenance systems, constraint satisfaction)
- Cognitive modelling (e.g. user models, memory models)
- Connectionist and PDP models
- Knowledge representation (e.g. logics for knowledge, belief and intention, nonmonotonic formalisms, complexity analysis, languages and systems for representing knowledge)
- Learning and knowledge acquisition
- Logic Programming (e.g. semantics, deductive databases, relationships to AI knowledge representation)
- Natural language (e.g. syntax, semantics, discourse, speech recognition and understanding, natural language front ends)
- Philosophical foundations
- Principles of AI applications (e.g. intelligent CAI, design, manufacturing, control)
- Qualitative reasoning and naive physics (e.g. temporal and spatial reasoning, reasoning under uncertainty, model-based reasoning, diagnosis)
- Robotics (e.g. kinematics, manipulators, navigation, sensors, control)
- Social, economic and legal implications
- Vision (e.g. colour, shape, stereo, motion, object recognition, active vision, model-based vision, vision architectures and hardware, biological modelling)

Timetable

Papers must be received by 10 December 1990, and must be 2500 to 5500 words in length. Authors will be notified of the program committee's decision on or before 20 March 1991.

For further information, including specific instructions for submitting a paper, contact one of the Program Committee CoChairs:

Prof. John Mylopoulos
Department of Computer Science
University of Toronto
Toronto, Ont. M5S 1A4, Canada
Tel: 1 (416) 978-5379
Fax: 1 (416) 978-1455
email: ijcai@cs.toronto.edu

Prof. Ray Reiter
Department of Computer Science
University of Toronto
Toronto, Ont. M5S 1A4, Canada
Tel: 1 (416) 978-5379
Fax: 1 (416) 978-1455
email: ijcai@cs.toronto.edu

Videotape Track:

Submission Requirements and Guidelines.

This track is reserved for displaying interesting research on applications to real-world problems arising in industrial, commercial, government, space and educational arenas. It is designed to demonstrate the current levels of usefulness of AI tools, techniques and methods. all video submissions will be peer reviewed.

Timetable

Submissions must be received by 10 December 1990. The notification of the decision will be made on or before 20 march 1991. For further information, including specific instructions of submitting a videotape, contact the Videotape track Chair:

Dr. Alain Rappaport
Neuron Data
444 High Street
Palo Alto, CA 94301, USA
Tel: 1 (415) 321-44488
Fax: 1 (415) 321-3728
email: atr@ml-ri.cmu.edu

IJCAI-91

n Artificial Intelligence

Panels

Individuals wishing to organize and chair a panel are invited to submit a suitable proposal. A panel allows three to five individuals to present their views or results on a common theme, issue. or question.

Panels should be both relevant and interesting to the AI community, and have a clearly specified topic that is narrow enough to be adequately addressed in a single session of slightly over one hour.

Timetable

Panel proposal should be submitted as soon as possible, but no later than 1 February 1991. Proposals will be reviewed as soon as they are received.

For further information, including specific instructions for submitting a proposal, contact the Panel Program Chair:

Dr. Peter F. Patel-Schneider
AT&T Bell Laboratories
600 Mountain Avenue
Murray Hill, New Jersey 07974 , USA
Tel: 1 (201) 582-3399
Fax: 1 (201) 582-5192
email: pfps@research.att.com

Tutorial Program

Proposals are invited form individuals wishing to offer a tutorial at IJCAI-91. Tutorial topics should be of interest to a substantial segment of the IJCAI audience. Proposals from a pair of presenters will be strongly favored over ones from a single individual.

Timetable

Proposals must be received by 4 Jan. 1991. Decisions about topics and speakers will be made by 22 Feb. 1991. Speakers should be prepared to submit completed course materials by 1 July 1991.

For further information, including specific instructions for submitting a proposal, contact the Tutorial Program Chair:

Dr. Martha Pollack
Artificial Intelligence Center
SRI International
333 Ravenswood Ave.
Menlo Park, CA 94025, USA
Tel: 1 (415) 859-2037
Fax: 1 (415) 326-5512
email: pollack@ai.sri.com

(Note: Indicate clearly on the first page that it is intended for Martha Pollack, Artificial Intelligence Center.)

Workshop Program

Gathering in an informal setting, workshop participants will have the opportunity to meet and discuss selected technical topics in an atmosphere which fosters the active exchange of ideas among researchers and practitioners. Members from all segments of the AI community are invited tn submit proposals for workshops they wish to organize.

To encourage interaction and a broad exchange of ideas, the workshops will be kept small, preferably under 35 participants. Attendance should be limited to active participants only. The format of workshop presentations will be determined by the organizers proposing the workshop, but ample time must be allotted for general discussion. Workshops can very in length. but most will last a half day or a full day.

Timetable

Proposals should be submitted as soon as possible, but no later than 21 December 1990. Proposal will be reviewed as they are receive and resources allocated as workshops are approved. Organizers will be notified of the committee's decision no later than 15 February 1991.

For further information, including specific instructions for submitting a proposal, contact the Workshop Program Chair:

Dr. Josph Katz
MITRE Corporation
MS-K318
Burlington Road, USA
Tel: 1 (617) 271-8899
Fax 1 (617) 271-2423
email: katz@mbunix.mitre.org

community . . . . Ironically, this brings us back to the central point of the naive physics manifesto . . . , one decade earlier.

It seems that Weld and de Kleer now have more respect for knowledge-oriented formalization methods but apparently no satisfactory method for formalizing qualitative reasoning has yet been found.

Selection: In choosing papers for the collection, Weld and de Kleer have attempted to be "completist", rather than "selective". I scanned my own bibliography on qualitative reasoning and every major paper on my list from before 1989 is in the collection. The editors have succeeded in including all major papers from their "school of thought", but have included few papers giving other views, apparently because few have been published in the AI literature. As seems typical in AI (unfortunately!), previous work in mathematics, such as the qualitative theory of dynamic systems, has only recently been identified as relevant. It is to the editors' credit that they have obtained and included Peter Struss's newly revised paper *Problems of interval-based qualitative reasoning*, which rigorously criticizes some of the assumptions of the field. For example, Struss argues that it is inconsistent with the history of science to derive a qualitative model from a quantitative model. In 1989, Sandewall criticized the field as based on the erroneous notion that one could not apply reasoning about inequalities directly to the traditional representation of differential equations, but his paper apparently appeared too late to be considered for inclusion in this collection.

Contents: The book contains 55 articles, divided into the following nine chapters:

1. Overview and motivation
2. Qualitative simulation
3. Mathematical aspects of qualitative reasoning
4. History-based simulation and temporal reasoning
5. Other styles of reasoning
6. Automating quantitative analysis
7. Multiple ontologies and automated modelling
8. Reasoning about shape and space
9. Causal explanation of behavior

Each chapter has a two- to six-page introduction by Weld and de Kleer. The authors represented most (by the number of papers for which they are listed as authors or co-authors) are as follows: Forbus 8 papers, de Kleer 5, Kuipers 3, Raiman 3, Sacks 3, and Weld 3.

Some of the highlights of the contents are now described.

Chapter l: The first chapter begins with Forbus's survey article from *Exploring artificial intelligence* (which is apparently a revised version of his *Annual review of computer science* paper, although this is never stated). This paper provides a survey of the field, but as a general introduction, it is perhaps marred by overemphasis on work done under Forbus's supervision at Urbana-Champaign. The next two papers are representative of the founding influences of de Kleer and Hayes on the research area. De Kleer's 1977 paper, *Multiple representations of knowledge in a mechanics problem-solver*, describes his NEWTON program, which uses quantitative techniques to resolve the ambiguities generated by envisioning (later known as qualitative simulation). Hayes's *Second naive physics manifesto* proposed encoding commonsense knowledge of the everyday world, as has already been mentioned, and also proposed the idea of "history", a piece of spacetime with natural spatial and temporal boundaries, as a primitive constituent for such descriptions.

Chapter 2: The second chapter is the longest and most crucial in the book. It gives foundational papers by de Kleer and Brown, Forbus, and Williams (updated version) from the *AI* journal special issue, which describe various qualitative versions of calculus and techniques for performing qualitative simulation. The chapter also includes papers presenting more recent developments in qualitative simulation theory. *Qualitative simulation* proceeds by following each qualitatively distinct path from each choice point, unless some constraint rules out a particular choice, until all possible behaviors have been generated. In *Qualitative simulation*, Kuipers proves that qualitative simulation does not miss possible behaviors, but does predict unrealizable behaviors. Separate papers address the intractable branching of qualitative simulation (Kuipers and Chui), reducing the number of choice points by additional constraints derived from the qualitative theory of dynamic systems (Lee and Kuipers; Struss), and extending qualitative simulation to include cases where discontinuous change occurs at choice points (Nishida and Doshita).

Chapters 3-5: The next three chapters describe the other research closely related to the central notion of qualitative simulation. The mathematical properties most often used for making the types of qualitative distinctions considered in qualitative reasoning are inequalities and orders of magnitudes. Chapter 3 explores the qualitative calculus (Struss; Williams), inequality reasoners (Simmons; Sacks), and order of magnitude reasoners (Raiman *et al*; Mavrovouniotis and Stephanopoulos). Chapter 4 presents William's research on incorporating Hayes's histories idea into qualitative simulation and includes representative papers on temporal reasoning by Allen, Vilain *et al*, and Dean and Boddy. Chapter 5 explores some of the styles of reasoning, other than qualitative simulation, that can be applied to a qualitative model. Papers by Weld (two), E. Davis, and Dague *et al* consider how a system's behaviour will change if its structure is changed, and one paper by Forbus gives a preliminary sketch of how a specific qualitative behavior can be identified from a qualitative model and observational data.

Chapters 6-9: The remaining chapters cover the "fringe areas" of qualitative reasoning. The subject of chapter 6, automating quantitative analysis, is not really part of qualitative reasoning, but people inevitably wonder about taking this alternative approach. The research reported here is based more rigorously on a quantitative model than is the previously discussed work on qualitative reasoning. Papers by Yip and by Sacks describe attempts to create qualitative summaries in the form of phase diagrams (on cartesian coordinates) of the most significant aspects of a quantitative model. Struss, in the paper included in Chapter 2, applies similar ideas to qualitative simulation.

Chapter 7 is a grab bag of nine articles, four of which deal with the problem of concurrently using several abstractions for the same physical system. Collins and others (two papers) discuss various models of the evaporation process in a refrigerator, Kuipers discusses distinct models for different time rates, and Hobbs identifies problems in relating separate models at different granularities. Of the remaining papers in the chapter, only Weld's *The use of aggregation in causal simulation* deals with automated modelling; he proposes a mechanism for abstracting simple loops into single operations. Automated modelling is identified by Weld and de Kleer as a daunting research problem.

The "qualitative kinematics" research described in Chapter 8 attempts to create qualitative techniques appropriate for reasoning about three-dimensional objects in space. Papers are included by Davis, Faltings, Forbus *et al*, Joskowicz (two papers), Gelsey, and Nielsen. Forbus *et al* claim that a purely qualitative approach to kinematic reasoning is underconstrained, and choose a combined quantitative/qualitative technique. Two interesting outside influences on the research on qualitative kinematics are (1) Reuleaux's early (1876) work on the theory of kinematics and (2) Lozano-Perez's work in robotics on configuration spaces.

Research in Chapter 9 tackles the difficult problem of causality, with little success. Rieger and Grinberg's 1977 "thermostat" paper provides a complex set of link types for specifying a causal model as a graph. Weld and de Kleer criticize this model in detail in their introduction, but this work has not been superceded. The three-paper controversy of de Kleer and Brown versus Iwasaki and Simon from *AI* journal is included, although these papers may have shed more heat than light on the matter [of causality] (Forbus, p. 26). It is interesting to see Weld and de Kleer state that de Kleer and Brown's paper fails to adequately address the issue "under dispute" (p.615). So little progress has been made on causality that de Kleer, in his IJCAI-87 *Computers and thought* speech, used the flush toilet as an example of what qualitative reasoning cannot handle, although the flush toilet was one of Rieger's examples in his own *Computers and thought* speech 12 years earlier.

Market: *Readings* is best suited as a reference work for researchers interested in qualitative reasoning, but the market for which Morgan Kaufmann is probably aiming is graduate-level survey classes in qualitative reasoning. The collection has been used by Jim Greer for a seminar class at the University of Saskatchewan. He reports that *Readings* was a good collection of important papers in the area. However, it was at times tough reading for graduate students because it lacks sufficient tutorial material to identify the principal ideas of qualitative reasoning and because the introductions to the chapters did not offer enough guidance to a reader new to the area. Perhaps what is missing is a tutorial paper that applies qualitative reasoning, as it is now understood, to a few simple applications. Such a paper would be particularly appropriate since Weld and de Kleer say "there is no single paper that summarizes our current understanding of qualitative simulation" (p. 84).

Index: One bonus with this collection is the indexes. I tested the name index by picking either numeric or alphabetic citations anywhere in the text of any paper and checking for the corresponding name in the index; in each case it was there. For example, a book by Christenson and Voxman is cited in a footnote on p. 270; both names appear in the name index, each with one entry, 270. The name index is less useful for researchers such as de Kleer and Forbus, who have entries for more than a hundred pages, including many consecutive pages. The two-level subject index is great!

Packaging: *Readings* is adequately packaged. The text is clearly printed, in spite of having been obtained by photo-reproducing existing papers. The collection is bound rather cheaply, like a conference proceedings, but the binding will probably last as long as book the is relevant. Flipping through the book may be a bit disconcerting because some of the papers are printed sideways. Nonetheless, buying this collection is much easier than tracking down all the papers.

References

Bobrow, D.G. (1984). Qualitative reasoning about physical systems: An introduction". *Artificial intelligence*, 24: 1-5.

Lenat, D.B., Prakash, M., and Shepherd, M. (1986). CYC: Using commonsense knowledge to overcome brittleness and knowledge acquisition bottlenecks". *AI Magazine*, 6(4): 65-85.

Sandewall, E. (1989). Combining logic and differential equations for describing real-world systems". In *Proceedings, First International Conference on Principles of Knowledge Representation and Reasoning (KR-89)*, Toronto, Ontario. 412-420.

Howard J. Hamilton is a Ph.D. candidate in the School of Computing Science at Simon Fraser University. His research interests include inductive inference and qualitative reasoning.

Probabilistic reasoning in intelligent systems: Networks of plausible inference

Judea Pearl
[UCLA]
San Mateo: Morgan Kaufmann Publishers, 1988, xix+552 pp
(The Morgan Kaufmann series in representation and reasoning) Hardbound, ISBN 0-934613-73-7, us $39.95

Reviewed by
Fahiem Bacchus
University of Waterloo

Over the past eight years, Pearl and his colleagues have done an impressive amount of work on applications of probability in AI. Pearl's book reports on the majority of this work, and its bulk (more than 500 pages) testifies to the quantity of work they have done in this area.

The book consists of ten chapters. The first two are introductory: Chapter 1 is an introduction to applications of probability in AI, and Chapter 2 is an introduction to Bayesian inference, the foundation of Pearl's approach to probabilities. The next six chapters comprise the bulk of the book, and contain its most significant contributions: network-based probability models. Chapter 9 contains a comparison of the Bayesian approach with other approaches to uncertainty management, notably the Dempster-Shafer theory. The final chapter contains Geffner and Pearl's epsilon-probability approach to non-monotonic reasoning.

Despite some attempts (e.g., the last chapter) to deal with other issues, this is essentially a book on network probability models, focusing in particular on Bayesian networks. Outside of publications in the scholarly journals, it contains the most detailed exposition of this kind of probabilistic model currently available. As I noted above, the middle six chapters are entirely devoted to an exposition of these models. We learn their formal definitions (Chapter 3); how they can be used for probabilistic reasoning (Chapter 4) and diagnosis applications (Chapter 5); how they can be extended to deal with decisions (Chapter 6), taxonomic hierarchies, continuous variables, and higher-order probabilities (Chapter 7); and how they can be constructed or learned from raw data (Chapter 8).

Because of its narrow focus, this book is not really suitable as a textbook, except perhaps in a rather specialized seminar course. I would consider it to be more of a reference book on Bayesian networks and their applications. It should be noted that by itself the book is not even broad enough to be used as a text in a course on the uses of probability in AI. There is a *lot* more to probability theory than the Bayesian approach or Bayesian networks! The book also reads much like a reference book, the presentation is not much better than that contained in some of the original journal articles. Nevertheless, there are many useful applications of Bayesian networks in areas such as expert systems that have to deal with uncertainty, and the approach has gained considerable attention in statistics as an aid to the statistician for conceptualizing decision problems. If you are interesting in these types of applications, this book will be indispensable.

My main criticisms of the book are twofold: (l) Pearl presents probability theory solely as the theory of Bayesian probabilities, where updates are always performed by conditioning (using Bayes's formula); and (2) the formalisms presented are essentially designed for probabilities attached to propositions, not to first-order assertions.

There is much more to probability theory than the Bayesian approach. For example, Pearl makes some rather weak arguments as to why beliefs should combine like frequencies (Section 1.4.1). The connection between frequency information (statistical information) and an agent's degrees of belief is a subtle and difficult problem. It is a problem that is not properly addressed by the Bayesian approach to probabilities, and it is a problem that is glossed over by Pearl.

The second problem becomes apparent in the last chapter, where he tries to attach probabilities to first-order assertions, making the ridiculous statement since classical logic does not possess an operator equivalent to the conditioning bar in probability, we are at liberty to modify the traditional treatment of the universal quantifier, whenever it contains a conditioning bar in its scope" (p. 475). Fortunately for us, he does not so abuse a standard concept. In the approach he presents, universal quantification is interpreted in the standard manner: as applying to all individuals. Unfortunately for him, this means that his approach leads to a rather questionable model that has many counterintuitive consequences, e.g., a strange and arbitrary division between evidence and background context that quickly falls into inconsistency if evidence is let into the background context. See Bacchus (1990) for a detailed exposition of how first-order assertions and probabilities interact, and the problems with Pearl's approach.

In sum, I would recommend this book as a reference text for those who would be interested in some of the very useful applications of Bayesian networks. However, if you want to learn something about probabilities in general and their possible applications to AI, this book is far from being complete.

Reference

Bacchus, Fahiem (1990) *Representing and reasoning with probabilistic Knowledge*. Cambridge, MA: The MIT Press [to appear].

Fahiem Bacchus is a member of faculty at the University of Waterloo. He is the author of the forthcoming book *Representing and reasoning with probabilistic Knowledge*.

Deduction systems in artificial intelligence
Karl Hans Äsius and Hans-Jürgen Bürckert (editors)
[IBM Stuttgart and Universität Kaiserslautern]
Chichester, England: Ellis Horwood, 1989, 238 pp
(Ellis Horwood series in artificial intelligence)
(Distributed in Canada by John Wiley and Sons Canada Ltd)
Hardbound, ISBN 0-7458-0409-8 and 0-470-21550-X, cdn $103.95

Reviewed by
Robert Mercer
University of Western Ontario

The AI, theorem-proving, and logic-programming communities have diverged in recent years. *Deduction systems in artificial intelligence* provided me with a reasonably up-to-date synopsis of the theorem-proving field, together with a snapshot of the intersection of it and the logic-programming community. It accomplished this task without getting sidetracked into the more esoteric areas of each field. As well, because of its European authorship, it includes material that is often missing in North American texts.

The goal of the book is, to quote the authors from the preface, "to give the reader an easily understandable yet comprehensive and up-to-date view of deduction systems". In a book of less than 240 pages, some of this grand plan must be sacrificed. But it gathers together a broad spectrum of topics (and a reasonably complete sample of the subject). Here is an outline of each of its five chapters.

Chapter 1, 'The history of deduction systems and some applications" (by J. H. Siekmann), is a generous serving of the history of reasoning, logic, and automated deduction. In just a few pages it threads together names such as Leibniz, Boole, Frege, Hilbert, Skolem, Gödel, Davis, Newell, Shaw, and Simon, as well as Robinson and many others, breathing a life into the topic that is discussed in the next 200 pages. It is a pleasant introduction not only to the book, but also provides easy access into the less-technical literature through a long list of references at the end of the chapter.

Chapter 2, "The foundations" (by N. Eisinger and H. J. Ohlbach), takes nearly half of the book. It is divided into four main sections. The first, an introduction to first-order predicate calculus, should be familiar to the intended readership. The next section contains descriptions of three deductive calculi: Gentzen calculi, the more familiar resolution, and theory resolution. The third section, which is half the chapter, discusses three forms of representation for deduction: tableaux, clause graphs, and matrices. The issue discussed, performance versus representation, which should be very familiar to computer scientists, is well presented. The last section is concerned with control issues within an automated deduction system.

Chapter 3, "The equality relation", is divided into four sections, each written by a different set of authors. Each section presents a different view of this important problem. Section 1 (by K.-H. Bläsius and H. J. Ohlbach) is a short introduction to the problem of incorporating the natural axiomatization of equality into an automated deduction system. Section 2 (by K.H. Bläsius) describes some general approaches to the equality problem, such as paramodulation, E-resolution, RUE-resolution, and equality graphs. Section 3 (by H.J. Bürckert) discusses various kinds of unification, including the obvious Robinson unification, as well as the lesser-known theory unification. This is an important section, as logic programming systems are beginning to incorporate this type of unification. Section 4 (by N. Eisinger and A. Nonnengart) provides a good introduction to term rewriting systems, and includes a lengthy discussion of the Knuth-Bendix procedure and some extensions.

Chapter 4, "Computational logic" (by H.-J. Bürckert), provides a fairly standard short introduction to logic programming, with which the intended readership should be quite familiar. However it also includes two short sections on sorts and types and feature types, which is usually not included in this type of short introduction.

Chapter 5, "Complete induction" (by D. Hutter), concludes the book with a discussion of computing inductive inferences from axiomatic theories.

The reader requires some mathematical sophistication. An appropriate readership would include graduate students and researchers in computer science who want an introduction to automated deduction systems. Except for the first two chapters, which comprise half the book, I found the material rather dense and at times lacking the "introductory" ingredient. However, like a tourist guide of a grand old city, the book indicates the historically important and presently trendy sites, possibly giving a not always immediately understandable introduction to each of them, and assumes that the visitor will spend many hours becoming acquainted with the various aspects of the city, usually with reference to other sources of information and by immersing him- or herself in the attractions.

The book is well-written. It is quite an undertaking to attempt to edit a book written by seven different authors. The editors have managed to keep an even tone throughout, although some areas were emphasized more than maybe necessary and some issues are completely missing. For example, the section on representation in Chapter 2 is half the chapter. While this may be an important topic for theorem proving, it seems to have less importance for those who may want to use a theorem prover. On the contrary, no mention that a sorted logic theorem prover was the first to solve Schubert's steamroller (a problem known to the AI community) is given. The only instance of a notational mismatch was the use of *suc* in one chapter and *succ* in another. The preface suggests that one can easily read chapters of the book independently of the other parts, except for the foundations chapter. I agree, and in addition because no special notation is used, one can easily read the last three chapters without having read the second chapter, if one is sufficiently grounded in logic.

For a couple of reasons I have never been fond of book titles of the form "Subject X for Audience Y", especially when the subject appears to be of interest to the audience yet the two are somewhat divergent. Firstly, without special care, the authors of the book (usually experts in subject X) write about what concerns those interested in subject X, rather than providing audience Y with a view of the subject from a viewpoint that is familiar to them — for example, using well-motivated examples from Y's problem domain. The present book unfortunately fits into this category, and for that reason will probably be ignored by much of the community (here I am thinking mainly of the scruffy camp). Others (the neats) will definitely find some enjoyment in seeing a formal treatment of problems that are of interest to the neat community. The second reason that I don't like this kind of title is that if the book is written in the manner in which this one has, it is useful to others outside community Y. This is especially true of the topics found in this book. I am quite sure that novices (or those wishing to expand their knowledge) in logic programming, software engineering (specification and verification), and databases (deductive databases) would find this book worth reading.

I found that the translation was nearly impeccable. Except for a few ill-chosen words that are easily reinterpreted, the only word that seemed inappropriate because of the introductory nature of the book (but my ignorance may be showing) was the word "junctor" instead of "logical operator". A careful reader may notice that two years transpired between the original German version of the book (in 1987) and the English translation. I think that it would have been appropriate to have included a preface to the translated version not only to warn the reader but also to give the editors a place to add some additional, more up-to-date references, remembering that the book is intended as an introduction.

Two of my pet peeves are found in this book; one concerns content, one form. In the second chapter one finds the notions that tautologies are valid statements (I have read "are" as "are the same as" rather than "are a subset of" so I may have misinterpreted the sentence) and that "inconsistent" is a synonym for "unsatisfiable". As well, the deduction theorem is presented as a semantic rather than as a syntactic notion. An introductory text is not the place to have such potentially confusing statements. Concerning form, references are given at the end of each chapter rather than collected at the end of the book. While this style lends itself to self-contained chapters, it makes locating references somewhat difficult. On the positive side however, the editors do include a page of general references at the end. The list is composed of the standard texts in logic and automated theorem proving.

The book suffers somewhat from sloppiness in form. Errors that are easily noticed (for example, an obvious error in a running title, missing column headings in tables) as well as less obvious (but somewhat numerous) typos makes one suspicious about possible errors in the more fragile and bountiful mathematical formulas. Many of the figures had inadequate accompanying explanation. And, in some of the lists of references, page numbers are not included.

To summarize: I believe this to be a good introduction to automated deduction systems for anyone who has some mathematical sophistication. This would include the intended readership of AI practitioners, as well as those of other computer science fields.

Robert Mercer is a member of the Department of Computer Science, University of Western Ontario. His research includes topics in default reasoning and linguistic pragmatics.

Automated generation of model-based knowledge-acquisition tools
Mark A. Musen
[Stanford University]
London: Pitman and San Mateo, CA: Morgan Kaufmann Publishers, 1989, xviii+2 93 pp (Research notes in artificial intelligence) Paperbound, ISBN 1-55860-090-6, us $24.95 (Distributed in Canada by John Wiley and Sons Canada Ltd)

Reviewed by
Stephen Regoczei
Trent University

To reduce the cost of knowledge acquisition, it was suggested that we put software tools in the hands of experts. This way they could do their own knowledge capture without having to deal with computer specialists who, typically, (bright as they may be), still have to be trained first in the use of the vocabulary and basic concepts of the application domain. This may seem like a good idea, but software tools, usable by experts without further technical assistance, have to be carefully tailor-made for the acquisition of a particular type of knowledge in a given domain area. Hence the cost of knowledge acquisition is now replaced by the expense of forging software tools.

The work reported in this book is based on an ingenious idea: to reduce the cost of toolmaking, let us build software that builds tools. This is certainly possible in some highly structured domains of expertise such as clinical trials for cancer therapies.

Musen's work is based on construing knowledge acquisition as modelling, and the possibility of creating skeletal models that can be further refined by adding more detail. The knowledge is to be represented as a special type of frame: namely, a form with blanks to be filled in. These forms can be extendible such as spreadsheets, or look more like the preprinted forms that we fill out at income tax time. These forms can be pictured as organized in a taxonomic hierarchy based upon generalization and specialization. Knowledge acquisition becomes equivalent to filling out forms. The skeletal outline is extended through what may casually be called "knowledge stuffing".

The practical implementation of this idea made use of ONCOCIN, an expert system for cancer treatments. Physicians wanted to have dozens of cancer-treatment plans, also called protocols, available in ONCOCIN's knowledge base, but the labour required for knowledge entry was prohibitive. In examining the problem, it became clear to Musen that the similarities among all the oncology protocols formed the basis for a general model. This model was used to drive a software tool called OPAL to generate and enter cancer protocols directly into ONCOCIN's knowledge base. In OPAL, however, the structure of the various graphical forms and their blanks can not be changed without reprogramming. A higher-level tool, called PROTEGE, was needed to edit the application-based assumptions in OPAL, and use different assumptions to generate other, OPAL-like software tools. Thus information entered into PROTEGE defines the skeletal structures, and PROTEGE generates families of OPAL-like software tools. In turn, these tools can be used by experts to directly produce knowledge bases for ONCOCIN-like advice systems. Thus PROTEGE is described as a 'tool to produce tools to produce tools".

Clearly, this work depends on the special structure of the application area. There are explicit conceptual models of the task domain. Communication difficulties are minimized by the well-harmonized, shared knowledge of the medical community, as well as by the fact that there is a one-to-one correspondence in this domain between words and concepts. But the basic ideas should be usable in other areas where skeletal-plan refinement is applicable.

The book, in addition to describing the author's recent research, also gives a broad, comprehensive review of recent knowledge acquisition work by others. Because Musen is writing for two communities, AI and medical, he gives careful definitions of the most important terms in both areas (e.g., describes some plausible reasons for distinguishing between "knowledge acquisition" and "knowledge engineering"). I found the book useful as a reference; I would recommended it both for the clarity of its writing and its imaginative extension of current knowledge acquisition technology.

Stephen Regoczei is an Associate Professor of Computer Studies at Trent University, Peterborough, Ontario. He is the author of several papers on manual knowledge acquisition techniques.

BRIEFLY NOTED

Foundations of cognitive science
Michael I. Posner (editor)
[University of Oregon]
Cambridge, MA: The MIT Press, 1990, xiv+888 pp
Hardbound, ISBN 0-262-16112-5, us $45.00

An invitation to cognitive science. Volume 1: Language
Daniel N. Osherson and Howard Lasnik (editors)
Cambridge, MA: The MIT Press, 1990, xix+273 pp Hardbound, ISBN 0-262-15035-2, usS37.50; paperbound, ISBN 0-262-65033-9, us $18.95

The MIT Press seems to be making a specialty of introductory collections on cognitive science. A couple of years ago, it published the first textbook in the area, *Cognitive science: An introduction*, edited by Neil Stillings and others. Now, two more collections have appeared. The first is a three-volume series, *Invitation to cognitive science*, edited by Daniel Osherson. The series is intended for students, or for researchers in other fields who want a nuts-and-bolts introduction to cognitive science. The first volume covers language from the perspective of theoretical linguistics, psychology, and philosophy. The next two volumes, to be published this year, will be entitled *Visual cognition and action* and *Thinking*. In contrast, Posner's collection (which is as big as Osherson's volumes put together), while introductory, seems to be intended for researchers rather than students, and emphasizes the coherence of the foundations and principles of cognitive science. — *G.H.*

BOOKS RECEIVED

Books listed below that are marked with a + will be reviewed in a future issue. Reviewers are still sought for those marked with a *.

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 below. All books received will be listed, but not all can be reviewed.

Readers who wish to review books for *Canadian AI* should write, outlining their qualifications, to the book review editor, Graeme Hirst, Department of Computer Science, University of Toronto, Toronto, Canada M5S IA4. Obviously, we cannot promise the availability of books in anyone's exact area of interest.

***Maureen Caudill and Charles Butler Naturally intelligent systems**
Cambridge, MA: The MIT Press / Bradford Books, 1990, ix+304 pp Hardbound, ISBN 0-262-03156-6, us $19.95

Common Lisp programming for artificial intelligence
Tony Hasemer and John Domingue
[Open University] Wokingham, England: Addison-Wesley, 1989, xiii+444 pp (International computer science series) Paperbound, ISBN 0-201-17579-7

Computer architecture: A quantitative approach
John L. Hennessy and David A. Patterson
[Stanford University and University of California, Berkeley, resp.] San Mateo, CA: Morgan Kaufmann Publishers, xxviii+748 pp (Distributed in Canada by John Wiley and Sons Canada Ltd) Hardbound, ISBN 1-55860-069-8, S65.95

Expert systems in data processing: Applications using IBM's KnowledgeTool (TM)
Joseph L. Hellerstein, David A. Klein, and Keith R. Milliken
[IBM T.J. Watson Research Center] Reading, MA: Addison-Wesley, 1990, xix+332pp Hardbound, ISBN 0-201-19540-2

Ethical conflicts in information and computer science, technology, and business
Donn B. Parker, Susan Swope, and Bruce N. Baker
[SRI International] Wellesley, MA: QED Information Sciences, 1990, 214 pp Paperbound, ISBN 0-89435-313-6, us $29.95

The VMS user's guide
James F. Peters III and Patrick J. Holmay
[Kansas State University and Analysts International Corp, resp.] Bedford, MA: The Digital Press, 1990, xxiii+304 pp Paperbound, ISBN 1-55558-014-9

+Developing and managing expert systems: Proven techniques for business and industry
David S. Prerau
[GTE Laboratories] Reading, MA: Addison-Wesley, 1990, xv+353 pp Hardbound, ISBN 0-201-13659-7

Advanced Prolog: Techniques and examples
Peter Ross
[University of Edinburgh] Wokingham, England: Addison-Wesley, 1989, x+294 pp (International series in logic programming) Hardbound, ISBN 0-201-17527-4

Working with analogical semantics: Disambiguation techniques in DLT
Victor Sadler
[BSO/Research, Utrecht]Dordrecht: Foris, 1989, 256 pp (Distributed language translation 5) Hardbound, ISBN 90-6765-429-9, us $60.00; Paperbound, ISBN 90-6765-428-0, us $33.00.

+Innovative applications of artificial intelligence
Herbert Schorr and Alain Rappaport (editors)
[USC Information Sciences Institute and Neuron Data, resp.] Menlo Park, CA: AAAI Press, 1989, xvi+363pp (Distributed by The MIT Press) Paperbound, ISBN 0-262-69137-X, us $19.95; Hardbound, ISBN 0-262-19294-2

Golden Common Lisp: A hands-on approach
David J. Steele
Singapore: Addison-Wesley, 1989, xxxii+543 pp Paperbound, ISBN 0-201-41653-0

World Watch

1.0 THEORETICAL ASPECTS

2 Bounds on the learning capacity of some multi-layer networks.
G.J.Mitchison, R.M.Durbin
(King's Coll. Res. Centre, Cambridge, UK).
Biol. Cybern. (West Germany). vol.60, no.5, p.345-56 (1989).

The authors obtain bounds for the capacity of some multi-layer networks of linear threshold units. In the case of a network having n inputs, a single layer of hidden units and an output layer of s units, where all the weights in the network are variable and $s \leq h \leq n$, the capacity m satisfies $2n \leq m \leq nt \log t$, where $t=1+h/s$. They consider in more detail the case where there is a single output that is a fixed Boolean function of the hidden units. In this case the upper bound is of order nh logh bus the argument which provided the lower bound of 2n no longer applies. However, by explicit computation in low dimensional cases they show that the capacity exceeds 2n but is substantially less than the upper bound. Finally, a learning algorithm for multi-layer networks with a single output unit is described. This greatly outperforms back propagation at the task of learning random vectors and provides further empirical evidence that the lower bound of 2n can be exceeded. (13 refs.)

3 The mind at AI: horseless carriage to clock [aims of AI].
W.C.Hill
(Microelectron. & Comput. Technol. Corp., Austin, TX, USA). AI Mag. (USA). vol.10, no.2, p.29-41 (Summer 1989).

Commentators on AI converge on two goals which they believe define the field: to obtain a better understanding of the mind by specifying computational models, and to construct computer systems that perform actions traditionally regarded as mental. It is argued that AI has a third, hidden, more basic aim; that the first two goals are special cases of the third: and that the actual technical substance of AI concerns only this more basic aim. This third aim is to establish computational based representational media, media in which human intellect can come to express itself with different clarity and force. This article articulates this proposal by showing how the intellectual activity known as AI can be likened in revealing ways to each of five familiar technologies. (11 refs.)

15 The dead end of symbolic AI and the connectionist approach.
M.F.Peschl
(Inst. fuer Angewandte u. Numerische Math., Tech. Univ., Wien, Austria).
Seventeenth Annual ACM Computer Science Conference, Louisville, KY, USA, 21-23 Feb. 1989 (New York, NY, USA: ACM Press 1989), p.427

This paper shows, using philosophical as well as logical arguments. the deficiency of the belief that universal intelligence could be based on the 'physical symbol systems hypothesis'. In spite of the (sometimes outstanding) results of orthodox AI, it seems that AI finds itself in a dead end. The thought experiment of the Chinese room shows best the problem of symbolic AI: obviously there exist more levels of understanding-it can be seen as manipulating symbols (which have in fact no meaning) or as a subsymbolic process of spreading (nervous) activations. There is shown in this paper why manipulating meaningless symbols are not a sufficient means for intelligent action.

251 Introduction: paradigms for machine learning.
J.G.Carbonell
(Sch. of Comput. Sci., Carnegie-Mellon Univ.. Pittsburgh, PA, USA).
Artif. Intell. (Netherlands), vol.40, no.1-3. p.1-10 (Sept. 1989).

A history of the development of machine learning is given, stressing its centrality to the field of AI. The four major paradigms in machine learning are then described; they are the inductive paradigm, the analytic paradigm, the genetic paradigm and the connectionist paradigm. Cross-paradigmatic observations are presented. (16 refs.)

262 A continued investigation into qualitative reasoning about shape and fit.
S.P.Carney
(Symbolics Inc., Cambridge, MA, USA), D.C.Brown.
(Al EDAM) Artif. Intell. Eng. Des. Anal. Manuf. *(UK)*, vol.3, no.2, p.85- 10 (1989).

Previous research on qualitative reasoning about shape and fit laid the foundations to determine whether two objects fit together. Continued investigation has refined the theory and has produced a functioning implementation. The paper describes extensions to the theory and the detail of the implementation. The reasoning process has been divided into five layers: grouping, topology, orientation, matching, and confirmation. The grouping layer clusters features such as cubes or cylinders into groups for each surface of an object. The topology layer recognizes patterns formed by the groups on each surface, and describes the pattern in terms of topological structures. The orientation layer selects promising surfaces from the two objects and attempts to align the two surfaces. It the orientation layer aligns the topological structures on the two surfaces, the matching layer tries to pair the features within the topological structures. The confirmation layer inspects paired features to determine whether the surfaces are compatible. If the surfaces are compatible, then the two objects qualitatively fit together. (21 refs.)

270 Artificial neural network on a SIMD architecture.
J.R. Brown, M.M. Garber, S.F. Venable
(Martin Marietta Electron. Syst., Orlando, FL, USA).
Proceedings. The 2nd Symposium on the Frontiers of Massively Parallel Computation (Cat. No.88CH2649-2), Fairfax, VA, USA, 10-12 Oct. 1988 (Washington, DC, USA: IEEE Comput. Soc. Press 1989), p.43-7

An implementation of a fully connected artificial neural network using the multilayered perceptron model is described. The neural network is implemented on a systolic array processor based on the Geometric Arithmetic Parallel Processor (GAPP) chip. Arrays of GAPP chips make up a single-instruction multiple-data (SIMD) class machine which has fine-grained connections and is fully programmable. Previous application areas of the GAPP system are image/signal processing, computer vision, and knowledge-based processing. The neural network is a relatively new processing model for the GAPP, but one that readily maps onto the architecture of the overall array processor. The proof-of-concept neural network is a multilayered perceptron model which uses the back-propagation learning paradigm. This initial network has fewer than 100 nodes in three layers and is trained to recognize letters of the alphabet. (9 refs.)

2655 RI: a logic for reasoning with inconsistency.
M.Kifer
(Dept. of Comput. Sci., State Univ. of New York, Stony Brook, NY, USA), E.L.Lozinskii.
Proceedings. Fourth Annual Symposium on Logic in Computer Science (Cat. No.89CH2753-2). Pacific Grove. CA. USA, 5-8 June 1989 (Washington. DC, USA: IEEE Comput. Soc. Press 1989), p.253-62

The authors present a logic, called RI (reasoning, with inconsistency), that treats any set of clauses, either consistent or not, in a uniform way. In this logic, consequences of a contradiction are not nearly as damaging as in the standard predicate calculus, and meaningful information can still be extracted from an inconsistent set of formulas. RI has a resolution-based sound and complete proof procedure. It is a much richer logic than the predicate calculus, and the latter can be imitated within RI in several different ways (depending on the intended meaning of the predicate calculus formulas). The authors also introduce a novel notion of epistemic entailment and show its importance for investigating inconsistency in the predicate calculus. (25 refs.)

2654 On the complexity of epistemic reasoning.
M.Y.Vardi
(IBM Almaden Res., San Jose, CA, USA).
Proceedings. Fourth Annual Symposium on Logic in Computer Science (Cat. No.89CH2753-2), Pacific Grove, CA, USA, 5-8 June 1989 (Washington, DC, USA: IEEE Comput. Soc. Press 1989), p.243-52

A study is made of the complexity of the decision problem for epistemic logics based on R. Montague's (1968) and R. Scott's (1970) semantics. The interest is in finding out how assumptions about the agents' reasoning power affect the complexity of reasoning about the agents' knowledge. A spectrum of assumptions is studied, and it is shown that the complexity of the logic under different assumptions is always in NP or PSPACE. The mental faculty that raises the complexity of the logic from NP to PSPACE is pinpointed. It is the ability to combine distinct items of knowledge. (27 refs.)

2663 1988 AFIT neural network research.
S.K.Rogers, M.Kabrisky
(USAF Inst. of Technol., Wright-Patterson AFB, OH, USA).
Proceedings of the IEEE 1989 National Aerospace and Electronics Conference NAECON 1989 (Cat. No.89CH2759-9), Dayton, OH, USA. 22-26 May 1989 (New York, NY, USA: IEEE 1989), p.688-94 vol.2

The authors provide a summary of recent research at the Air Force Institute of Technology (AFIT) in the area of neural networks. Specifically, AFIT research in the areas of error drive learning algorithm acceleration, speech recognition, target classification, time series prediction, and optical and VLSI implementations is presented. AFIT has developed and tested an algorithm which fits a curve to the error surface of a backpropagation surface to find more quickly a good set of weights. AFIT has investigated the use of momentum and second-order error-driven learning algorithms. A speech recognition system was developed using, a combination of Kohonen nets and dynamic time warping. In the area of target classification, the accelerated backpropagation nets have been successfully applied. Research into predicting chaotic time series was performed. Three optical neural network architectures were designed and tested, and a VLSI implementation was investigated.

2.0 SYSTEMS AND TECHNIQUES

20 KIMS — a knowledge-based computer vision system for production line inspection.
C.A.Ntuen, E.H.Park
(Dept. of Indust. Eng., North Carolina, A&T State Univ., Greensboro. NC, USA), J.H.Kim.
Comput. Ind. Eng. (UK) vol.16, no.4, p.491-508 11989).

Discusses a generic expert system prototyped for image analysis and interpretation tasks. The system christened KIMS is a knowledgebased image management system which oversees the entire process of image processing, segmentation, feature extraction, knowledge representation along with an expandable capability for image understanding task. The authors discuss the architecture of KIMS, its modeling environment, and demonstrate its usefulness with photo-to-image inspection as applied to a manufacturing line. (13 refs.)

32 An integrated system for design of mechanisms by an export system — DOMES.
B.Yang, U.Datta. P.Datseris, Y.Wu
(Rhode Island Univ., Kingston, Rl. USA).
(AI EDAM) Arrif. Intell. Eng. Des. Anal. Manuf. (UK). vol.3. no. 1, p.5370 11989).

Methodologies have been developed and implemented in LISP and OPS-5 languages which address type synthesis of mechanisms. Graph theory and separation of structure from function concepts have been integrated into an expert system called DOMES (design of mechanism by an expert system) to effectively implement the following three activities: enumeration of all non-isomorphic labelled graphs: identification of those graphs which satisfy structural constraints: and sketching of mechanisms corresponding to a given graph. Developed theories and algorithms are applied to a robot gripper design and a variable-stroke piston engine design. The results from these two applications indicate that the automated techniques effectively identify all previously obtained solutions via manual techniques. Additional solutions are also identified and several errors of the manual process are detected. The developed methodologies and software appear to perform a complete and unbiased search of all possible candidate designs and are not prone to the errors of the manual process. Other important features of DOMES are: it can learn and reason, by analogy, about a new design problem based on its experience of the problems previously solved by the system: it has the capability to incrementally expand its knowledge base of rejection criteria by convening into LISP code information obtained through a query-based interactive session with a human designer: and it can select the set of rejection criteria relevant to a design problem from its knowledge base of rejection criteria. These procedures could become a powerful tool for design engineers, especially at the conceptual stage of design. (19 refs.)

44 Justifying the cost of expert system development.
A.M.O'Connell
(Texas Instrum. Inc., Dallas. TX, USA).
ESD/SMI Expert Systems Proceedings, Detroit, MI, USA, 12-14 April, 1988 (Detroit. MI, USA: Eng. Soc. Detroit 1988), p.279-94

This paper discusses the process of cost justification. and the estimates of costs and benefits which are the starting, point for project justification. Checklists of typical expert system costs and benefits ere presented as estimating guidelines. Actual cost and benefit information about several expert systems developed by Texas Instruments is presented. Three methods for performing cost/benefit analysis are explained, with comments on their relative merits for different situations. An example expert system development project is analyzed using all three methods. (5 refs.)

38 Automating knowledge acquisition — a survey of systems and approaches.
L.O'Connor. D.R.Panridge. C.P.Dolan, N.Ebeid. N.H.Goddard
(Artificial Intelligence Center, Hughes Res. Labs., Calabasas, CA. USA). ESD/SMI Expert Systems Proceedings, Detroit, Ml, USA, 12-14 April, 1988 (Detroit. Ml, USA: Eng. Soc. Detroit 1988), p.181-201

Knowledge acquisition has long been a bottleneck in knowledge based expert systems development. Consequently, techniques for automating the knowledge acquisition process have been sought since the inception of this field. Recently. there has been a virtual explosion in the number and types of systems designed specifically for automating knowledge acquisition — in universities, industrial research centers, and AI software houses This paper summarizes the results of a recent state-of-the-art technology assessment by the Hughes AI Center. An analysis of the knowledge acquisition process is presented, which leads to a hierarchic view of knowledge acquisition featuring four levels of abstraction. The paper features a survey of research and development activity in systems for automating knowledge acquisition. An extensive bibliography is included. (80 refs.)

285 An architecture for integrating expert systems.
N.M.Fraser
(Dept. of Manage. Sci., Waterloo Univ., Ont., Canada),
K.W.Hipel, D.M.Kilgour, M.D.McNeese, D.E.Snyder.
Decis. Support Syst. (Netherlands), vol.5, no.3, p.263-76 (Sept. 1989).

A conflict resolver is developed for mediating disputes arising among competing expert systems (ESs). The conflict resolver constitutes the key component of an overall decision support system designed for controlling a complex system or supporting the decision-making of a human operator within the system. When competing ESs make conflicting recommendations for controlling the system, the conflict resolver uses a Delphi-like mediation to achieve a compromise solution. The mathematical framework for the comprehensive conflict resolution system is defined formally, and it is proven that the conflict resolver always reaches a consistent recommendation. (21 refs.)

286 The Canadian expert systems market: the fourth wave.
EDP In-Depth Rep. (USA), vol.18, no.10, p.1-15 (Oct. 1989).

Players in the expert systems market are battling to overcome both product misconceptions as well as challenges associated with developing practical applications using leading edge technology.

Presently in the awareness and acceptance stage of its product life-cycle, expert system technology is beginning to be recognized by both end-users and management as a weapon providing strategic advantage. The onus, however, is on vendors and developers to push the concept of expert system technology into the mass market. Data collected for this report suggests that expert system technology will be part of mainstream computing by the mid-to-late 1990s.

291 Expert system problem selection: a domain characteristics approach.
H.K.Jain
(Sch. of Bus. Adm., Wisconsin Univ., Milwaukee, WI, USA), A.R.Chaturvedi. Inf. Manage. (Netherlands), vol.17, no.5, p.245-54 (Dec. 1989).

The first generation of commercial expert systems based on AIO technology are available in the market place. But in the available literature, one can find hardly any material on expert system problem selection. A number of popular and successful expert systems are analyzed. Domain-dependent and domain-independent problem characteristics have been identified, based on the analysis. To test the contention that these characteristics significantly contribute to the success of expert systems, a questionnaire survey involving a number of expert system developers was conducted. Based on this, a domain characteristic approach for expert system problem selection is presented. (16 refs.)

294 Personal Consultant plus [expert system tool].
Microcomput. Civ. Eng. (USA), vol.4, no.3, p.243-4 (Sept. 1989).

Developed by Texas Instruments Inc., and introduced to the market in 1987, Personal Consultant Plus (or PC Plus) is a member of the family of Personal Consultant expert system tools for building expert systems on personal computers. The low-end member of the family is the Personal Consultant Easy (or Pc Easy). PC Easy and PC Plus have a lot in common. The article reviews only the additional features and capabilities of PC Plus Release 3.0.

2675 PC-based expert system shells: some desirable end less desirable characteristics.
R.G.Vedder
(BCIS Dept., North Texas Univ., Denton, TX, USA). Expert Syst. (UK), vol.6, no.1, p.28-42 (Feb. 1989). [received: 29 Aug 1989]

The marketplace offers a great variety of PC-based expert system shells. More than ever, ES developers need a clear idea of what characteristics they require in an ES toolkit. The author discusses a recent project which used different shells to build versions of a common application. Evaluating this work illuminates some of the important attributes one should look for. The project specifically involved the examination of five different PC-based shells—Exsys, 1st-Class, Guru, Personal Consultant Easy and Personal Consultant Plus. Teams of graduate students with knowledge engineering experience used these tools to build a common prototype application. What they discovered provides a valuable look at the usefulness of many shell characteristics. (3 refs.)

2677 Issues in the verification of knowledge in rule-based systems.
D.L.Nazareth
(Sch. of Bus. Admin., Wisconsin Univ., Milwaukee, WI, USA). Int. J. Man-Mach. Stud (UK). vol.30, no.3, p.255-71 (March 1989). [received: 26 Sep 1989]

As expert system technology spreads. the need for verification of system knowledge assumes greater importance. The paper addresses the issues involved in demonstrating a rule-based system to be free from error. A holistic perspective is adopted. when in sources, manifestations and effects of errors are identified. A general taxonomy is created, and the implications for system performance and development outlined. Existing strategies for knowledge verification are surveyed, their applicability assessed, and some directions for systematic verification suggested. (65 refs.)

2680 Measurement issues in knowledge engineering.
L.Adelman
(Dept. of Inf. Syst. & Syst. Eng., George Mason Univ., Fairfax, VA, USA). IEEE Trans. Syst. Man Cybern. (USA), vol.19. no.3, p.483-8 (May-June 1989).

There are five sources (for determinants) of knowledge base quality: domain experts, knowledge engineers, knowledge representation schemes, knowledge elicitation methods, and problem domains. The knowledge base for many expert systems is developed for a problem domain using one domain expert, one knowledge engineer, one knowledge representation scheme, and one elicitation method. Since there is minimal research demonstrating that the possible variation in each of these sources does not significantly affect the quality of the knowledge base, the generalizability (or validity) of such systems in real-world settings is questionable. Consequently. research is needed to assess the extent to which system validity is affected by these sources of variability. Toward this end, the results of reanalyzing the data from an experiment varying domain experts, knowledge engineers, and elicitation methods when developing a multiattributed representation scheme for combat readiness are presented. No significant effects were obtained for elicitation method or knowledge engineer. (129 refs.)

2697 Artificial neural network simulation on Cray systems.
D.A.Olsen, S.C.Ahalt, R.S.Barga, G.L.Wilcox.
Cray Channels (USA). vol. 11, no.2, p.20-3 (Summer 1989).

Researchers are designing artificial neural networks (ANNs) that Simulate nerve cells, or neurons, their interconnections, and patterns of interaction. They are using Cray computer systems to simulate, study, and develop numerous artificial neural network algorithms. Cray systems lend themselves naturally to ANN simulations because most of the calculations performed during the simulations are vector oriented, and the simulations are characterized by fine-grained parallelism, which is well-supported by pipelined vector processors. The authors look at distributed artificial neural network simulation and large-memory artificial neural network simulations. They conclude that although no one can predict the direction and success of future ANN technologies, in the short term two developments seem very likely. The first is the software 'hybridization' of ANNs with conventional software systems and with expert systems. The second likely development is the hardware implementation of ANNs in analog VLSI circuitry. (3 refs.)

2708 An expert ECG acquisition and analysis system.
A.Perkusich, G.S.Deep, M.L.B.Perkusich, M.L.Varani
(Dept. of Electr. Eng., Univ. Fed. of Paraiba, Brazil).
Conference Record. IEEE Instrumentation and Measurement Technology Conference (Cat. No.89CH2707-8), Washington. DC, USA, 25-27 April 1989 (New York, NY, USA: IEEE 1989), p.1849

The design aspects and actual implementation phase of an expert system as an aid for the analysis and diagnostics of cardiac infirmities is presented. An IBM-PC compatible microcomputer containing the necessary interface cards for the acquisition of ECG (electrocardiograph) signals is used. The system provides electronic switching of the ECG lead configurations. A data and control flow diagram is used for the analysis of the software of the proposed system. Particular attention is given to signal acquisition, signal processing data, data management, the knowledge base inference engine, and the man-machine interface. (11 refs.)

2705 Of brittleness and bottlenecks: challenges in the creation of pattern-recognition and expert-systems models.
M.A.Musen, J.van der Lei
(Dept. of Med. Inf., Erasmus Univ., Rotterdam, Netherlands).
Pattern Recognition and Artificial Intelligence. Towards an Integration. Proceedings of an International Workshop, Amsterdam, Netherlands. 18-20 May 1988 (Amsterdam, Netherlands: North-Holland 1988), p.336-52

Concerns the problem of eliciting useful and complete knowledge about a given application task from experts — the knowledge-acquisition bottleneck. The problem of brittleness in expert systems – their marked degradation in reasoning performance when confronted with unusual or atypical cases — is also addressed. Most expert systems like pattern-recognition systems, are concerned with the classification of entities. The process of building knowledge bases for expert systems that perform classification, like the process of developing statistical classifiers, requires that someone who is familiar with the application area both determine the relevant types of objects and identify their germane observable features. In the case of both expert systems and pattern-recognition systems, developers must create models of the application area. Although the pattern-recognition literature does not generally mention concepts such as 'knowledge acquisition' or 'brittleness'. these problems are also important in the construction of statistical models. This paper shows how builders of expert and pattern-recognition systems face many of the same challenges, and discusses ways in which the two research communities can learn from each other's experiences. (31 refs.)

3.0 APPLICATIONS

90 Planning for expert system implementation in the 'real world'.
D.L.Smith
(Adv. Manuf. Eng., Ford Motor Co., Milan, MI, USA).
ESD/SMI Expert Systems Proceedings. Detroit, MI, USA. 12-14 April, 1988 (Detroit. MI. USA: Eng. Soc. Detroit 1988), p.15-26

Provides practical guidance to individuals interested in creating and implementing AI based expert systems in the real world, as opposed to laboratory situations. Combining knowledge, gained through observation of and involvement in the implementation of various forms of advanced technology in the manufacturing environment, with the experiences of others, the author has identified a number of the pitfalls that can adversely affect efforts to implement technology. Through a discussion of some of the pitfalls that are relevant to the implementation of expert systems, he provides some useful hints for avoiding these pitfalls while creating and implementing expert systems. (13 refs.)

113 A real time expert control strategy for blood gas management in neonates under ventilation treatment
L Zhang, R G Cameron
(Dept of Control Eng. Bradford Univ UK)
IEE Colloquium on Exploiting the Knowledge Base: Applications of Rule Based Control (Digest no.89), London, UK, 1 June 1989 (London, UK IEE 1989), p.4/1-7

In the therapy of premature infants with respiratory distress syndrome, their arterial oxygen partial pressure is routinely controlled by adjusting the concentration of oxygen (F_IO_2) in the air they inspire. For more severe cases where the patient is unable to breath unaided, assisted ventilation is introduced to maintain the oxygen supply and carbon dioxide elimination. Multiple ventilators controlled parameters, including peak inspiratory pressure positive end-expiratory pressure, ventilation frequency , inspiratory to expiratory duration ratio may be altered skilfully to optimize gas exchange. Various rules have been derived by paediatritions based on their clinical experiences. In the paper the authors discuss an expert ventilator manager combined with a realtime generalised predictive controller. The objective of the system is to control both blood oxygen and carbon dioxide pressures within their prespecified ranges by adjusting F_I0_2 and four ventilator factors. (12 refs.)

180 Methodology for implimenting artificial intelligence systems in Ada.
B. Abrams, T. Doran
(Dept of Software Syst Grumman Aircraft Syst. Div Bethpage, NY USA)
Proceedings of Sixth National Conference on ADA Technology Washington, DC, USA 14-18 March 1988 (Springfield, VA, USA U. S. Dept. Commerce 1988), p 602-8

The languages most commonly used to prototype artficial inteligence (AI) systems are very different from the languages used to program avionics software. The paper proposes a methodology to develop arificial intelligence software for an avionics application dispite this mismatch. The methodology involves prototyping in an AI language, then reimpliminting in Ada using a library of AI function. The pilot project used to verify this methodology is a maintenance diagnostic system prototyped in Lisp and reimplimented in Ada (13 refs.)

183 AI in CAD: a coupled expert system.
H. Adeli, Y. Paek
(Dept. of Civil Eng., Ohio State Univ., Columbus, OH, USA).
ESD/SMI Expert Systems Proceedings, Detroit, MI, USA,
12-14 April, 1988 (Detroit MI, USA Eng. Soc. Detroit 1988),
p 51-62

A knowledge-based expert is presented for detailed design of steel framed buildings, called STEELEX. STEELEX is a coupled system in which symbolic processing is coupled with numerical processing.In order to perform svmbolic and numerical processing, a structural defsign language (SDL) was developed in INTERLISP environment. SDL contains the inference engine, the explanation facility, the debugging facility, and redesign management. The complex body of knowledge needed for detailed design of a fractionated into smaller and manageable knowledge sources which are organized into a hierarchy of cooperating conceptual specialists. STEELEX presents the final detailed design of the steel structure including the connections. Thus, it can be also be used for manufacturing of the structure. (9 refs.)

184 Expert system support in engineering design.
O.K. Helferich, S.J. Young
(Dialog Syst. Div., A.T. Kearney Inc., East Lansing, MI. USA). ESD/SMI Expert Systems Proceedings, Detroit, MI, USA, 12-14 April1988 (Detroit. MI. USA: Eng. Soc. Detroit 1988). p.63-74

Expert systems continue to evolve as a viable means to obtain productivity increasing in industry. The paper covers a packaging design application. The development of an expert system which will assist the US Army and later industrial companies in the area of packaging design. Specifically, this expert system assist packaging engineers in the evaluation of packaging requirements and selection of the packaging solution for high value items that require cushion packages. The expert system is discussed in terms of the problem domain,solution, the potential benefits of applying the solution to the problem at hand, the knowledge engineering tools and techniques used in the development of the solution, and the lessons learned so far. Technical optics reviewed include knowlede acquisition, repr esentation and maintenance, as will as system integration issues.

196 Selection of appropriate solvents by an expert system based on databased technology.
A. Farny, H.Pfeiffer
(Inst. fur Inf., Tech. Univ. Munchen, West Germany). F. Bischofsberger, M.J. Hampe. Computer Application in the chemical Industry. Papers of a European Symposium, Erlangen, West Germany, 23-36 April 1989 (Weinheim, West Germany: VCH Verlagsgesellschaft 1989), p.517-24

It is frequently necessary to select an appropriate solvent for a desired separation process. A great amount of information about the physical data of the pure substances and their mixtures involved, about their chemical reactivity and about economical and ecological frame conditions can easily be stored in a database. The system for the selection of solvents, called SOLVENT, supports both a substrate-oriented retrieval of solvents according to the selection criteria specified by the user for a number of properties, and a problem-oriented selection of appropriate solvents for three types of application: azeotropic distillation, extractive distillation and liquid extraction. Rule systems, based of the stored solvent properties, are used for the case of the problem-oriented selection of solvents to specify the decision criteria. The result is a list of solvents which can be considered for a given separation problem. The estimation about the suitability of a solvent is made by means of evaluation functions and individual weighting factors for specific criteria.

204 An expert system for the conceptual design of bridges.
J. C. Miles, C.J. Moore
(Sch. of Eng., University of Wales, Cardiff, UK).
Artificial Intelligence Techniques and Application for Civil and Structural Enqineers, London. UK. 19-21 Sept. 1989 (Edinburqh. UK: CivilComp Press 1989). p. 171-6

The design of any structure can be said to consist of two stages: the first being the conceptual design stage, inn which the overall form of the structure is decided upon. The second stage consists of a more detailed structural anaylsis, during which calculations are carried out to verify the conceptual design choice, and determinecomponent dimensions. Bridges are a prime example of the type of structure where conceptual design is of importance, and this paper describes the development of an expert system for use in the conceptual design of a bridge. The information has been obtained from 'expert' engineers, witha wealth of experience in bridge design, and this paper discusses the methods used to elicit the expert knowledge from these engineers. In addition, the work described has thrown light on the subject of conceptual design itself, an area poorly documented to date. It is shown that the use of an expert system would provide a representation of the way in which the conceptual design process is carried out, as well as providing a useful advisory facility for the inexperienced designer. (20 refs.)

205 An expert system for the preliminary design of timber roofs.
D.R.Thompson
(Dept. of Civil Eng.. Brighton Polytech., UK), T. Tomski, S.W. Ellacott, P. Kuczora.
Arificial Intelligence Techniques and Applications for Civil and Structural Engineers. London, UK. 19-21 Sept. 1989 (Edinburgh, UK: CivilComp Press 1989), p.187-96

An expert system to design timber roofs is being produced as part of a project to investigate the process of engineering design. The operations involved in design can be broken down into synthesis and analysis. Analysis can be performed bv known and understood methods so the main problem of designing can be identified as that of synthesis. It has been found that most commercially available expert system shells are too restrictive for the development of an engineering design expert system. In order to impliment a system that more closely models human thinking, a new shell is being developed in tandem with the timber roof expert system and another concerned with heat exchanger design. This shell is being constructed specifically for engineering design problems. The features in this shell are those that used by engineers in their reasoning, such as hypothesis proposal and modification, non-monotonicity as well as the more well established techniques of numeric reasoning, inherence and the relational modelling of knowledge. These features are illustrated with reference to the timber roof problem. (10 refs.)

209 An expert system for offshore structure inspection and maintenance.
A.Langdon, K.Ahmad
(Dept. of Math., Surrey Univ., Guildford, UK). P.A.Frieze. Anificial Intelligence Techniques and Applications for Civil and Structural Engineers, London, UK, 19-21 Sept. 1989 (Edinburah, UK: CivilComp Press1989), p.221-32

The inspection and maintanance of offshore structures is carried out by experts who have considerable experience, and are able to bring to bear a large body of knowledge, for solving a variety of problems.

However, the full extent of knowledge required to completely rationally execute inspection and maintenance of offshore structures is beyond the scope of one individual without the assistance of a comprehensive system with access to the necessary databases and algorithms. The PLAIM (Platform Lifetime Assessment through Analysis, Inspection and Maintanance) project has been established to provide such a systems. 'Knowledge-based' expert system methods and techniques, particularlv those related to the acquisition of experiential knowledge and the representation of the knowledge for solving problems, have been used in the PLAIM project..The major delivarables of this project include: the documentation of informal and qualitative knowleddge together with a prototype expert system for eliciting details of defects in structures. empirically assessing the severity of the defect and recommending a series of remedial actions. (13 refs.)

215 The application of artificial intelligence to civil and structural engineering. A bibliography.
B.H.V.Topping, B.Kumar
(Dept. of Civil Eng., Heriot-Watt Univ., Riccarton, UK). Arificial Intelligence Techniques and Applications for Civil and Structural Engineers. London, UK, 19-21 Sept. 1989 (Edinburph. UK: CivilComp Press1989), p.285-303

Presents a bibliooraphy of papers concerned with the application of artificial intelligence techniques to civil and structural engineering. The topics covered include: expert systems shells, logic programming, interfaces and knowledge elicitation; uncertainty and fuzzy logic systems; graphics; water engineering; environmental engineering; geotechnical engineering; materials engineering; pavement engineering; transport and traffic engineering; surveying and road layout; architecture, layout, space and building layout; building design and assessment: construction engineering; construction planning, management and control: structural analysis and design; constraints, codes, rules and design: database systems: bridge engineering; structural loading offshore engineering; dynamics and earthquake engineering; structural optimization; finite element analysis, modelling and idealisation; and education. (436 refs.)

234 Automatic speech recognition — can it improve the man-machine interface in medical expert systems?
J.A.Landau. K.H.Norwich, S.J.Evans
(Dept. of Physiol., Toronto Univ.. Ont.,Canada). Int. J. Bio-Med Comput. (Netherlands), vol.24, no.2, p.111-17 (July 1989).

Computer-aided medical diagnosis has existed for two decades, but has not yet attained widespread acceptance among physicians. It is proposed that automatic speech recognition may be a significant factor in the eventual acceptance of the technolooy by the medical profession. The current state-of-the-art ot automatic speech recognition is briefly surveyed, and problems with the technology are discussed. A potential natural language interface with DIAG. an expert system for aiding in dermatologic diagnosis, is described. A system that has been developed for accepting input of body parts in freestyle format is presented as a prototype for a natural language interface with an automatic speech recognition device. (21 refs.)

235 An expert system for the selection of reusable program modules.
R.Zimbel, P.Weber
(Forschungszentrum Inf., Karlsruhe, West Germany). Inform. Forsch. Entwickl. (West Germany). vol.4, no.4, p.174-92 (1989). In German.

A software development environment based on the idea of reusing program modules has to support the tasks of module design, module selection and module integration. The expert system SESMOD (Selection Expert System for Modules) aims at the second task: it supports the selection of modules deposited in a library. The user specifies his requirements for the desired module in a dialogue. SESMOD compares those requirements with the properties of all modules in the library, and selects and rates the suitable ones. SESMOD is based on a taxonomy for program modules which contains declarations of domain-specific notions together with their relationships, The taxonomy and the descriptions of the modules in the Iibray, are interpreted by the expert system shell SESE. which is tailored to catalogue selection problems. (26 refs.)

238 Speech recognition bibliography.
N.B.Lerner
(Texas Univ., Austin, TX, USA).
SIGOIS Bull. (USA), vol.10, no.3, p.1-13 (Julv 1989).

A bibliography is presented covering various types and aspects of speech recognition. Computerized speech recognition, new developments, applications and AI involvement are some of the areas discussed. The author provides an abstract for each book or article in the bibliography. (22 refs.)

330 Proceedings of the First Annual Conference on Innovative Applications of Artificial Intelligence.
Menlo Park, CA, USA: American Assoc. Artificial Intelligence (1989), xiii + 185 pp.

Conference held at: Stanford, CA, USA. Date 28-30 March 1989. The following topics were dealt with: financial trading;

logistics management system; planning the discharging and loading of container ships; space shuttle mission control; statistical precess control; computer operation automation; personal finance planning system; financial analysis of automobile dealers; coal-mining emergency control advisory system; PepPro peptide synthesis expert system; Enhanced System Monitor; plastic food packaging advisor for marketing resin; television advertising campaign planning; Ford Motor Company direct lavor management system; inventory reduction with CAN BUILD; MACPLAN for airlift planning; intelligent training system for space shuttle flight controllers; harnessing detailed assembly process knowledge; cooling system design assistant; hydraulic circuit design assistant; intelligent banking system; IC design critic; music composition by emotional computation; naval battle management decision aid; AMTS automated money transfer service; legal consultation system; credit authorization system for American Express; ManTall rescue operations assistant; and CHARLEY for diagnostics of manufacturing equipment.

341 Review of expert systems in auditing.
D.E.O'Leary, P.R.Watkins
(Univ. of Southern California, Los Angeles, CA. USA).
ExpertSyst. *Rev. (USA)* vol.2, no.1-2, p.3-22 (Spring-Summer 1989). [received: 05 Dec. 1989]

The current state of expert systems and decision support systems in auditing is examined. The authors Iook at completed or prototype expert systems and decision systems in both external and internal auditing, including special areas of focus such as EDP auditing, and governmental auditing. They do not discuss or try to differentiate between expert systems and decision support systems. Both types of systems support audit decision making and, thus, both are included. The authors discuss the audit environment and proceed to review, respectively, audit-based expert systems in EDP auditing, external auditing: academic systems, external auditing: commercial systems, government auditing and internal auditing. Limitations of auditing- based expert systems are discussed, followed by a discussion of sources for publication and presentation of information relating to expert systems. (101 refs.)

342 Accounting expert systems: a comprehensive, annotated bibliography.
C.E.Brown
(Oregon State Univ., Newport, OR. USA).
Expert Syst. *Rev. (USA)*. vol.2, no.1-2, p.23-129 (Spring-Summer, 1989). [received: 05 Dec 1989]

A bibliography is presented whose objective is to provide researchers with a comprehensive annotated reference list of work on expert systems for accounting. Because expert systems is a rapidly expanding and relatively new field, traditional methods of searching the literature are not always sufficient. Articles on similar subjects by different authors frequently contain bibliographies with little commonality in the sources cited. A comprehensive collection of both the traditional and the nontraditional sources of information clearly is needed. The comprehensive annotated bibliography of articles on expert systems applied to accounting-related problems is intended to meet that need and includes 478 entries (476 refs.)

351 *PFPS*—personal financial planning system.
K.W.Kindle, R.S.Cann, M.R.Craig
(Chase Lincoln First Bank. Rochester, NY. USA), T.J.Martin.
Proceedings of the First Annual Conference on Innovative Applications of Artificial Intelligence, Stanford. CA, USA. 28-30 March 1989 (Menlo Park, CA, USA: American Assoc. Artificial Intelligence 1989), p.38-14

PFPS is an expert system developed over the last five years at Chase Lincoln First Bank, NA. to provide objective affordable expert financial advice to individuals with household incomes ranging from $25000 to $150000, and up. PFPS is an integrated personal financial planning system encompassing the following areas of financial expertise: investment planning; debt planning retirement savings and settlement of retirement plans; education and other children's goal funding; life insurance planning disability insurance planning: budget recommendations: income tax planning and savings for achievement of miscellaneous major financial goals (purchase of house, extended travel, etc). PFPS has supported the personal financial planning service of Chase Lincoln First Bank in upstate New York since late 1987. The PFPS system components are implemented as an embedded system on an IBM 4300 series mainframe under VM/CMS and on a Symbolics 3600 series LISP processor. The Symbolics system is connected to the IBM mainframe in a master/slave relationship using an application-level protocol defined on top of RS232C. All of the inferencing and planning is done on the Symbolics 36xx using a system architecture based on a blackboard frame work, an object-oriented data base, and a goal directed generate-and-test search paradigm with extensive search-space pruning. (5 refs.)

352 *ANALYST an* advisor for financial analysis of automobile dealerships
M.A.Hutson
(Gen. Motors Acceptance Corp., Detroit, MI, USA).
Proceedings of the First Annual Conference on Innovative Applications of Artificial Intelligence, Stanford, CA. USA. 28-30 March 1989 (Menlo Park. CA. USA: American Assoc. Artificial Intelligence 1989), p.46-52

ANALYST is deployed in 230 GMAC domestic branch offices. There, each year, 400 credit analysts perform nearly 20000 financial analyses of automobile dealerships. Using ANALYST they will save an estimated 55000 man-hours. The focus of this knowledge-based system from the very beginning was its end user—a credit analyst. Driven by singular end user requirements, developers extended the native Al tool explanation facilities, built an easy to use interface and completed the integration of engineering workstations in GMAC's IBM computing environment.

356 Expert systems end their applications in production and operations management.
Y.P.Gupta
(Dept. of Manage., Louisville Univ., KY, USA). D.C.W.Chin.
Comput. Oper. Res. (UK), vol.16. no.6, p.567-82 (1989).

Examines the concepts of expert systems and reviews the literature pertinent to the areas of application in production management such as scheduling, layout planning and tools used in applying these concepts such as stimulation and optimization. It has been concluded that due to increasing automation requirements in manufacturing and the Intensified organizational competitive environment, the decision making process itself will have to be automated. It has been suggested that expert systems could be a vehicle in achieving this. As expert systems become more readily available, the management issues of safety, validity and reliability will become more crucial. The compensation on transferring of knowledge is unclear. Also, the fear of replacement of white collar jobs is growing. (135 refs.)

364 A planning expert system for part setup and workholding.
P.J.Englert
(AT&T Bell Labs., Whippany, NJ, USA). P.K.Wright. *Adv.* Manuf. Eng. (UK). vol. no.5, p.249-63 (Oct. 1989).

Describes an expert system that has been built for pan Setup and workholding in small-batch manufacturing. Knowledge engineering with experienced machinists has uncovered distinct setup patterns and tradeoffs that characterize how these craftsmen plan machine setup procedures and clamp parts. An interesting observation is that some of the design tradeoffs used by the craftsmen are related to quantitative physics-based analyses, whereas others are qualitatively based on experience. An expert system framework is described that incorporated these two types of knowledge and the physical description of clamping components. Case studies and hardware are described. The research impacts on computer aided design methods and expediting process planning for small-batch manufacturing. (9 refs.)

379 Packaging Advisor: an expert system for rigid plastic food package design.
A.S.Topolski, D.K.Reece
(E.I. du Pont de Nemours & Co. Inc., Wilmington, DE, USA). Proceedings of the First Annual Conference on Innovative Applications of Artificial Intelligence, Stanford. CA. USA. 28-30 March 1989 (Menlo Park, CA. USA: American Assoc. Artificial Intelligence 1989), p.72-7

In 1987, the Du Pont Company entered the market for barrier resins , which are used in the fabrication of plastic food containers. The company was experiencing difficulty establishing a position against incumbent competitors. A new and technically superior product was about to be introduced, and a way was needed to induce customers to invest in qualifying the new material. The solution was Packaging Advisor, an expert system for rigid food container design. Deployed in February of 1988, Packaging Advisor automates the design process, providing its customer, the package designer, with information on alternative material, the quantities of these materials needed to meet performance specifications, and estimates of material costs. Packaging Advisor was used in place of traditional marketing communications techniques to inform customers and field sales staff about new and existing products. Customer response, expressed in both words and order, has been enthusiastic.

382 A pedagogical computer-aided design expert system for compensator development.
G.B.Lamont
(Electr. & Comput. Eng., Air Force Inst. of Technol., Dayton. OH, USA). CoED *(USA). vol.9* no.4. p.17-28 (Oct.-Dec. 1989).

Contemporary computer-aided design (CAD) programs for control systems education incorporate a variety of functional tools but can be enhanced through integration of expert systems. Expert systems for the teaching of control system design, however, can incorporate 'design process models' explicitly. To study the various elements of such a process, a a specific expert system has been generated to guide the control engineering student through lead-lag compensator design. Conventional and Al programming concepts are combined to generate a generic 'design associate' for this application. The expert system called TOTAL-EASE is built on top of ICECAP-PC (interactive control engineering computer analysis package for personal computers) although interfacing to a different computer-aided control systems design (CACSD) package is possible. The author discusses the various developmental aspects of an educational expert system. The components of a CACSD expert system organization including human-machine interfacing, search processing and knowledge representation selections are analyzed for application to control engineering education. (31 refs.)

417 An expert system in chemical synthesis written in APL2/PC.
J.Rojas, P.Rodriguez, M.Alfonseca
(IBM Madrid Sci. Center, Spain). J.LBurrgos
APL Quote Quad (USA) vol.19, no.4, p.299-303 (Aug. 1989). (APL 89: APL as a Tool of Thought, New York, NY, USA, 7-10 Aug. 1989).

In 1988. the Scientific Center of IBM in Madrid started a joint project with the Institute Quimic de Sarria. The objective was the construction of an expert system that would be able to perform the completely automated design of the quasi-optimal synthesis paths to obtain certain products. The expert system in chemical synthesis has been programmed successfully, using APL2. This first version of the system did not require the logic auxiliary processor, and could make use of a simplified internal representation of molecules. The synthesizer can be extended in the future to other families of chemical products without great difficulty. In particular, the graphics editor already supports a very general set of products (much larger than the synthesizer) and would not need to b extended. (6 refs.)

434 An expert system for cogeneration energy system selection.
G.C.Birur
(Jet Propulsion Lab., California Inst. of Technol., Pasadena. CA. USA). R.Lee.
Proceedings of the 23rd Intersociety Energy Conversion Engineering Conference, Denver. CO. USA, 31 July-5 Aug. 1988 (New York, NY, USA: ASME 1988). p.129-34 vol.4

To help engineers inexperienced in cogeneration energy system selection, a demonstration expert system has been developed. It demonstrates how a short list of candidates can be composed and later evaluates in detail for the final selection. The expert system assists engineers by asking them questions about: the application; electrical and thermal loads; fuel price; and electrical rates. It uses the answers to the questions along with its knowledge base to make recommendations on the candidates that have the best potential for the final selection. The expert system knowledge based is derived from extensive experience gained during the selection and implementation of cogeneration energy systems for naval bases in the last ten years. The user will obtain the detailed perfomance of these candidates using a cogeneration analysis model. In this way, a new engineer can function like an expert designer and save considerable effort by not having to evaluate of a large list of potential candidates. Using one example site, the recommendations made by the expert system were compared with those made by an earlier cogeneration evaluation study that used the Civil Engineering Laboratory cogeneration analysis program. The comparison shows that the recommendations made by the expert system are in good agreement with those made in earlier detailed cogeneration studies. (5 refs.)

467 Building envelope analysis and design system.
P. Fazio, C. Bedard, K. Gowri
(Centre for Building Studies & Siricon, Concordia Univ., Montreal, Que. Canada).
Artificial Intelligence in Design. Proceedings of the Fourth International Conference on the Applications of Artificial Intelligence in Engineering, Cambridge, UK, 10-13 July 1989 (Southampton, UK: Comput. Mech. Publications 1989), p. 133-42

The selection of building envelope components is often subject to a number of conflicting performances requirements depending on climatic conditions, user requirements combinations of available materials and constructional types have to be considered for achieving economy and overall successful performance. Knowledge-based system techniques can be effectively used to solve some of these difficulties as shown in the present implementation. A knowledge base of constructional types and material properties is first integrated with building code specifications for establishing the quantitative performance requirements of a design context. A control strategy for selectively exploring the design decision space has been developed for generating and evaluating the alternatives. The paper discusses and reviews the building envelope design process, the knowledge-based system approach to solve this design problem. and details of a prototype system implementation. (14 refs.)

479 Real time expert system in biochemical process control.
A. Halme
(Lab. of Autom., Fac. of Inf. Technol., Helsinki Univ. of Technol., Espoo, Finland), M.N.Karim, J.Makela, M.Pokkinen, M. Kauppinen..
Proceedings of the ISMM International Symposium. Computer Applications in Design, Simulation and Analysis, Honolulu, HI, USA,1-3 Feb. 1988 (Anaheim, CA, USA: ACTA Press 1989), p. 30-4

A real-time expert system is being developed for monitoring and analysis, diagnosis of faults and malfunctions, and online optimization of batch fermentation processes. The main purpose of the expert system is to assist the operators in running and improving fermentation conditions with minimum supervision. The expert system is interfaced with a computer-based data acquisition and control system which gives the basic online information from the process. In addition to that the operator can give offline data which can include physical attributes of the precess (color, and general appearance of the fermentation broth) and microscopic states of the microorganisms. Computational information like estimates of variables can also be included. The export system has a knowledge base which includes typical fermentation data (characteristic curves) of commonly run fermentations, culture data, operational data, calibration of sensors and rulesets which conduct inferences. The main functions of the system are related to its fault detection capabilities for making intelligent decisions about the abnormalities in fermentation experiments. (2 refs.)

2727 An artificial intelligence program: management lessons learned.
R.M.Boan
(Artificial Intelligence Div., Wright-Patterson AFB, OH, USA).
Proceedings of the IEEE 1989 National Aerospace and Electronics Conference NAECON 1989 (Cat. No.89CH2759-9). Dayton. OH, USA, 22-26 May 1989 (New York, NY, USA: IEEE 1989), p.1374-5 vol.3

The US Air Force Logistics Command (AFLC) has taken a low-cost, low-risk, high-payback approach to the development of artificial intelligence (Al) applications which can improve effectiveness, efficiency, and productivity. This is accomplished with centralized management and decentralized development. Using this approach 15 systems have been fielded and over 250 are under development. These applications vary from electronic system diagnostics to nondestructive inspection. It is found that the truly pertinent issues in developing and fielding Al systems are managerial rather than technical in nature. Some of the pertinent managerial issues are discussed in the form of lessons learned .

2733 Knowledge-based systems for commercial line insurance. An exciting application of AI-techniques in the service field.
S.Akselsen, G.Hartvigsen, P.W.Richardsen
(Dept. of Comput. Sci., Tromso Univ., Norway).
AI Commun. (Netherlands), vol.2, no.2, p.9B-109 (June 1989).

Deals with the concept of knowledge-based systems (KBS) in the insurance business. The use of artificial intelligence (Al) techniques in the service sector implies many different applications. Activities thus far have been concentrated around the industrial and scientific sectors. Introduction of knowledge-based systems as tools in the insurance business is a fairly new approach. The insurance business includes applications for sale/purchasing, underwriting, claims adjusting. reserving and auditing. Focusing on the commercial line insurance, this paper gives a brief overview of research activities in the USA, Europe and Scandinavia, which includes both insurance and artificial intelligence. This paper demonstrates the applicability of knowledge-based systems on current problems in commercial line insurance. In closing. an ongoing project is presented with an example illustrating the practical use of such a system. (33 refs.)

2736 The scope of artificial intelligence literature: a review of publication outlets and information sources.
C.E.Brown
(Oregon State Univ., Corvallis, OR, USA).
Expert Syst. Rev. (USA) vol.1, no.4, p.30-7 (Sept. 1988).
[received: 16 Aug 1989]

The list of journals, magazines and newsletters was compiled to identify sources of information on expert systems and artificial intelligence as they relate to business and accounting.

2761 Computer-aided medicine: present and future issues of liability.
H.Mortimer.
Comput./Lsw J. (USA). vol.9 no.2, p.177-203 (Spring 1989).
[received: 09 Aug 19891

Liability in computer-aided medicine is addressed. The author focuses on expert systems currently used in diagnostics as well as future uses for computers in medicine. Presently, hospitals we expert systems to aid physicians and other medical personnel in the diagnosis and monitoring of patients. It is anticipated that expert systems for home diagnostic use also will become a reality in the near future. Bas d on public policy and fairness rationales she argues that manufacturers of mass-produced computer systems should be held to a strict liability standard for defectively produced systems. A negligence standard should govern when a system is custom-made or when a physician's services are used in conjunction with the system. (125 ref .)

2768 The SENEX system: a microcomputer-based expert system built by oncologists for breast cancer management.
J.L.Renaud-Salis. F.Bonichon, M.Durand, A.Avril. C.Lagarde, J.P.Serre. P.Mendiboure
(Fondation Bergonie, Bordeaux, France).
Proceedings AIME 87. European Conference on Artificial Intelligence in Medicine. Marseilles, France, 31 Aug.-3 Sept. 1987 (Berlin, West Germany: Springer-Verlag 1987), p.5470

The major general goal of the SENEX project is to demonstrate the feasibility of cancer experts developing efficient consultation systems that can run on affordable microcomputers. This could help to manage clinical trials and disseminate state of the art cancer treatment protocols issued from cancer clinical research. The first specific goal of the project was to develop a prototype system applied to breast cancer treatment (the SENEX system) with the aid of a knowledge engineering package (Personal Consultant Plus, Texas Instruments). The paper reports on the successive phases of the knowledge base development, the current system status and the problems encountered. The SENEX system has reached the stage of a research prototype. It embeds approximately 400 production rules attached to 40 frames and performs at the expert-level for in-protocols patients. Its control structures are modeled from the ONCOCIN system. (9 refs.)

2785 Pattern recognition and artificial intelligence in molecular biology.
G.Harauz
(Dept. of Molecular Biol. & Genetics. Guelph Univ., Ont., Canada).
Pattern Recognition and Artificial Intelligence. Towards an Integration. Proceedings of an International Workshop, Amsterdam, Netherlands, 1 B-20 May 1988 (Amsterdam, Netherlands: North-Holland 1988). p.437-47

Molecular biology can be defined as the study of the chemical and physical structures and functions of biological macromolecules, of which there are literally millions of different kinds. Advances in experimental methodologies have led to the elucidation of the three-dimensional configurations of numerous proteins, and to the determination of primary sequences of many more. Such work is inherently computation and information intensive. Statistical pattern recognition and machine intelligence techniques are thus emerging as powerful tools to assist in three-dimensional structure determination, and in unravelling the evolutionary and functional relationships between different macromolecules. An exciting goal of this research, not yet realized, is the ability to design molecules with specified electrical, catalytic. or pharmaceutical properties. (86 refs.)

2797 Principles of developing expert systems of traditional Chinese medicine.
Qin Dulie
(Dept. of Med. Inf. Res.. Capital Inst. of Med., Beijing, China).
Expert Systems and Decision Support in Medicine. 33rd Annual Meeting of the GMDS EFMI Special Topic Meeting. Peter L. Reichenz Memorial Conference. Hannover, West Germany, 26-29 Sept. 1988 (Berlin, West Germany: Springer-Verlag 1988). p.86-93

The author begins with an introduction to traditional Chinese medicine (TCM), discussing the development of TCM in China and TCM as a unique medical science. The motivation for the development of TCM expert systems includes the development of systems such as MYCIN and the desire of well-known TCM veterans to store their expertise in an electronic brain. Over 160 TCM ES have been developed is research supported by Beijing city government, 50 of these have passed formal evaluation. A table lists the diseases covered by them and the names of 9 TCM ES are given noting where they were developed. The principles of developing TCM ES are outlined. In particular knowledge representation is discussed and some variants of production rules used in TCM ES are described. The process of knowledge acquisition from experts is also discussed. Some temporal aspects of TCM ES. particularly the need to master the stages of diseases development and to master clinical combinations of the states of illnesses are discussed. Finally the limited software environment in China and the difficulties associated with input and output of Chinese characters are briefly mentioned. (11 refs.)

2827 HardSys—applying expert system techniques to electromagnetic hardening.

J.LoVetri, S.Abu-Hakima. A.S.Podgorski
(Nat. Res. Council, Ottawa, Ont.. Canada). G.l.Costache.
IEEE 1989 National Symposium on Electromagnetic Compatibility (Cat. No.89CH2736-7), Denver. CO. USA. 23-25 May 1989 (New York, NY, USA: IEEE 1989). p.383-5

The application of expert system techniques to the electromagnetic hardening domain is briefly described with regard to a prototype named HardSys, written in Prolog and containing a subset of the domain knowledge. In this system the domain knowledge is partitioned into four electromagnetic hardening considerations: the ambient field; the shielding effectiveness the system susceptibility and the probability of failure. Each of these hardening requirements is implemented as a separate adviser module. Each module can be used independently and each can use information produced by any of the other modules if that information is available and/or required. Such is the case. for example, with the probability of failure module which requires the knowledge deduced from the ambient field, the shielding effectiveness, and the equipment susceptibility advisers. Preliminary results are reported to be encouraging. (11 refs.)

2868 Preference control: a language feature for Ada applications.

T.Elrad
(Illinois Inst. of Technol., Chicago, IL, USA), F.Maymir-Ducharfne.
Proceedings of AIDA-87. Third Annual Conference on Artificial Intelligence and Ada, Fairfax, VA, USA, 14-15 Oct. 1987 (Fairfax, VA, USA: George Mason Univ. 1987). 14 pp.

One of the goals of artificial intelligence is to simulate the human decision making process. Ada's concurrency and its associated language features controlling nondeterminism in real time systems allow a partial simulation of this decision making process. A more complete specification of the human thought process requires new language features that will allow the programmer more expressive and dynamic programming power during decision making. The addition of the dynamic preference control feature to Ada gives programmers the power to go beyond making decisions solely considering Boolean conditions and intertask communication; it allows the assignment of different preferences to different decision alternatives within the select statement and these preferences can be assigned dynamically to better adapt to real-time changes. Artificial intelligence comes one step closer to reality with the enhancement of the Ada language to support dynamic preference control by improving the simulation of the decision making process. (18 refs.)

2844 DICODE—an expert system for designing distillation column controls.

G.Birky. T.McAvoy
(Dept. of Chem. Eng., Maryland Univ., College Park. MD. USA), B.Tyreus.
Proceedings of the 1989 American Control Conference (Cat. No.89CH2772-2), Pittsburgh, PA, USA. 21-23 June 1989 (Green Valley, AZ. USA: American Autom. Control Council 1989). p.95-100 vol. 1

An expert system for designing control systems for distillation columns, called DICODE (distillation column design expert), has been developed. It contains approximately 300 rules and it has gone through extensive prototyping. The authors focus on the distillation control aspects of the expert system. Initially, in developing DICODE the design methodology of P. Buckley (1985) was used and then that of G. Shinskey (1988) was added. Examples are presented that show that the two design methodologies produce results that are diametrically opposed to one another in a number of cases. An attempt is made to explain how such differences can arise. (12 refs.)

"Based on the information provided in the abstracts, the references provided have been selected by the secretariat of the NRC Associate Committee on Artificial Intelligence 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, United Kingdom. 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, N.J. 08855-1331, U.S.A."

Persons wishing to obtain copies of references cited should contact their nearest technical library or the Canada Institute for Scientific and Technical Information, National Research Council, Building M-55, Montreal Road, Ottawa, Ontario, CDA, K2A 052, (tel.: (613) 993-1585, Telex 053-03115). For on-line ordering, CAN/OLE users may use the CAN/DOC command. Envoy users type "COMPOSE CISTI."

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 you 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 ($35.00 Cdn.).

❑ I am a student (subtract $10.00 Cdn.)
and/or
❑ I am a member of CIPS (subtract $10.00 Cdn.)

❑ Please send me the following back issues of Canadian Artifical Intelligence ($10.00 Cdn. each including postage, $15.00 Cdn. for issue #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 ($20.00 Cdn.). ‡

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 fowards your order.

CSCSI/SCEIO Conference Proceedings

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

❑ # of copies ___ Edmonton, 1988 (Within Canada: members – $35.00 Cdn.; non-members – $40.00 Cdn. Add $5.00 Cdn. for postage. Mail to CIPS.) (Outside Canada: $35.00 US. Postage within U.S.: $2.00 US for the first copy, $.75 US for each additional copy. Postage outside U.S.: $3.00 US for the first copy, $2.00 US for each additional copy. Within California: add 7% sales tax. Mail to: Morgan Kaufmann Publishers, Order Fulfillment Center, P.O. Box 50490, Palo Alto, CA 94303, USA)

❑ # of copies ___ Montréal, 1986 ($30.00 Cdn. Add $5.00 Cdn. postage for within Canada, $7.00 Cdn. for outside Canada. Mail to CIPS.)

❑ # of copies ___ Saskatoon, 1982 ($25.00 Cdn. Add $5.00 Cdn. postage for within Canada, $7.00 Cdn. for outside Canada. Mail to CIPS.)

❑ # of copies ___ Victoria, 1980 ($25.00 Cdn. Add $5.00 Cdn. postage for within Canada, $7.00 Cdn. for outside Canada. Mail to CIPS.)

Name ______________________________

Mailing Address ______________________________

____

Canadian Artificial Intelligence

Intelligence Artificielle au Canada

National Research Council Canada | Conseil national de recherches Canada

Computational Intelligence | Intelligence informatique

Editors | Directeurs scientifiques
N. Cercone / G. McCalla

Computational Intelligence, the official journal of the Canadian Society for Computational Studies of Intelligence, is a quarterly journal first published in 1985 by the National Research Council of Canada. It contains high-quality theoretical and experimental research papers in computational (artificial) intelligence, by encouraging contributions from the following fields: knowledge representation; natural language understanding; computational vision; applications of artificial intelligence; logic programming; theorem proving; learning; cognitive science; problem solving and planning; languages and tools for artificial intelligence; speech understanding; game playing; philosophical implications; and foundations of artificial intelligence. Three special issues were published in 1988: "Taking Issue: an inquiry into Computer Understanding"(February), "AI in France"(May), and "Planning"(November). The Journal is international in content and distribution and is quickly becoming one of the leading AI journals in the world

La revue Intelligence informatique, organe officiel de la Société canadienne pour l'étude de l'intelligence par ordinateur, est publiée quatre fois par an par le Conceil national de recherches du Canada depuis 1985. Elle renferme des articles de qualité dans le domaine de l'intelligence artificielle, tant au niveau théorique qu'expérimental, et encourage la publication de communications dans les domaines suivants: la représentation des connaissances, la compréhension des langages naturels, la vision computationnelle, les applications de l'intelligence artificielle, la programmation logique, la démonstration de théorèmes, l'apprentissage, la science cognitive, la résolution et la planification de problèmes, les langages et les outils de l'intelligence artificielle, la compréhension de la parole, les jeux, la portée philosophique et les fondements de l'intelligence artificielle. Trois numéros spéciaux ont été publiés en 1988: Forum : "An Inquiry into Computer Understanding" (février), "AI in France" (mai) et "Planning" (novembre). La revue est internationale tant par l'origine des auteurs qui y contribuent que par sa distribution et est rapidement en voie de devenir une publication de pointe dans le domaine de l'IA.

Computational Intelligence is published quarterly. Subscription rates for 1990 are $48.00 for Individuals and $120.00 for Institutions.

La revue Intelligence informatique est publiée 4 fois par an. En 1990, l'abonnement individuel coûte 48$ et l'abonnement collectif 120$.

For further information contact

Subscription Information
Research Journals, M-55
National Research Council of Canada
Ottawa, Ontario, Canada
K1A OR6

Pour tout information complémentaire, communiquez avec

Abonnement
Revues scientifiques, M-55
Conseil national de recherches Canada
Ottawa (Ontario)
K1A 0R6

Canada